

System Design Methods for Simultaneous Optimal Control of Combustion Instabilities and Efficiency

Final Report
ONR Contract N00014-99-1-0752
September 30, 2002

Principal Investigators:

William T. Baumann
William R. Saunders
Uri Vandsburger

*College of Engineering
Virginia Polytechnic Institute and State University
Blacksburg, VA 24061*

DISTRIBUTION STATEMENT A:
Approved for Public Release -
Distribution Unlimited

20031010 096

REPORT DOCUMENTATION PAGE

Form Approved
OMB NO. 0704-0188

Public Reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comment regarding this burden estimates or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188,) Washington, DC 20503.

1. AGENCY USE ONLY (Leave Blank)		2. REPORT DATE September 30, 2002	3. REPORT TYPE AND DATES COVERED Final Report: 1 May 1999 to 30 Sept 2002
4. TITLE AND SUBTITLE System Design Methods for Simultaneous Optimal Control of Combustion Instabilities and Efficiency		5. FUNDING NUMBERS G N00014-99-1-0752	
6. AUTHOR(S) W. T. Baumann, W.R. Saunders, U. Vandsburger, A. Greenwood, M. Vaudrey, M. Carson		7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) VIRGINIA POLYTECHNIC INSTITUTE AND STATE UNIVERSITY	
8. PERFORMING ORGANIZATION REPORT NUMBER VACCG-ONR-1		9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES) <u>SPONSORING AGENCY</u> OFFICE OF NAVAL RESEARCH Ballston Centre Tower One 800 North Quincy Street Arlington, VA 22217-5600	
10. SPONSORING / MONITORING AGENCY REPORT NUMBER SPONSOR: ONR 251 MONITOR: N66020		11. SUPPLEMENTARY NOTES	
12 a. DISTRIBUTION / AVAILABILITY STATEMENT Distribution Unlimited		12 b. DISTRIBUTION CODE	
13. ABSTRACT (Maximum 200 words) The work under this contract primarily involved the design, analysis and testing of control algorithms aimed at the suppression of thermoacoustic instabilities and the optimization of combustor performance. Three different types of algorithms were considered in this work: Pattern search, explicit gradient, and least-mean-square based feedback, which we have designated filtered-E. The algorithms are listed in order of increasing speed and increasing amount of a priori information required. Both pattern search and explicit gradient algorithms are useful for the optimization of combustion performance as well as the suppression of thermoacoustic instabilities. Filtered-E is intended solely for the fast suppression of thermoacoustic instabilities. All of the algorithms were successful in suppressing thermoacoustic instabilities in an experimental combustor. In addition, we provide an analysis of the mechanism for achieving control using on-off actuators pulsed subharmonically, validate the analysis with experimental results and propose a variable-subharmonic controller. Control systems using on-off actuation can be adaptively tuned using pattern search or explicit gradient algorithms with little modification. The filtered-E algorithm has also been effective when applied to on-off actuation, even though the implicit gradient is not quite correct.			
14. SUBJECT TERMS Thermoacoustic instabilities, Pulsed control, Optimization algorithms, Control algorithms			15. NUMBER OF PAGES
			16. PRICE CODE
17. SECURITY CLASSIFICATION OR REPORT UNCLASSIFIED	18. SECURITY CLASSIFICATION ON THIS PAGE UNCLASSIFIED	19. SECURITY CLASSIFICATION OF ABSTRACT UNCLASSIFIED	20. LIMITATION OF ABSTRACT UL

Table of Contents

1	Introduction	1
2	Pattern Search Algorithms	4
2.1	Introduction	4
2.2	Theory	4
2.3	Implementation and Robustness Issues	6
2.4	Experimental Results	8
2.5	References	22
3	Explicit Gradient Algorithms	23
3.1	Introduction	23
3.2	Theory	23
3.3	Implementation and Robustness Issues	24
3.4	Experimental Results	25
3.5	References	37
4	Filtered-E Based Algorithms	38
4.1	Introduction	38
4.2	Adaptive Feedback Control	39
4.3	Analysis of Adaptive Feedback Applied to Self-Excited Systems	41
4.4	Simulation, Experimental Results, and Actuator Authority	46
4.5	Stability and Operating Constraints of Adaptive LMS Feedback Control	54
4.6	Practical Applications and Considerations	61
4.7	Simulation	65
4.8	References	69
5	Pulsed Control	72

5.1 Introduction	72
5.2 Analysis	73
5.3 Controller Implementation	79
5.4 Experimental Results	81
5.5 Conclusion	89
5.6 References	90
6 Conclusions	91

1 Introduction

1.1 Overview

The work under this contract primarily involved the design, analysis and testing of control algorithms aimed at the suppression of thermoacoustic instabilities and the optimization of combustor performance. Although these two problems are similar in that they can be cast as optimization problems, they have significant time scale differences: Thermoacoustic instabilities must be suppressed quickly, on the order of a second, whereas combustor performance optimization can take place on the timescale of minutes. In both cases, very little is known about low-order, control-oriented modeling of these processes. This necessitated the use of simple adaptive algorithms that require very little a-priori information about the system.

Three different types of algorithms were considered in this work: Pattern search, explicit gradient, and least-mean-squares (LMS) based feedback, which we have designated filtered-E. The algorithms are listed in order of increasing speed and increasing amount of a priori information required. Both pattern search and explicit gradient algorithms are useful for both the optimization of combustion performance and the suppression of thermoacoustic instabilities. Filtered-E is intended solely for the fast suppression of thermoacoustic instabilities.

In the pattern search category, this work investigates both the Hooke and Jeeves, and the Rosenbrock algorithms. These algorithms can operate with a controller of any structure and systematically perturb the free parameters until further improvement is not possible. The advantage of pattern search algorithms is that they are capable of searching very rough performance surfaces and require almost no a priori information. In addition, they should never allow the controller to remain in regions of the parameter space that degrade performance below the uncontrolled case. The major disadvantage is that they can be slower than other algorithms since they make no assumptions about the space they are searching. Hence, pattern search is ideally suited to slow optimization of combustion performance. In addition, we have found that it performs quite well for suppressing instabilities.

With the explicit gradient algorithms, three variations were considered: a simple time-averaged gradient (TAG), a gradient with linesearch, and a conjugate gradient method. These algorithms all explicitly compute a gradient of the performance with respect to the free parameters. For smooth performance surfaces, they are faster than the pattern search algorithms. Their main disadvantage is the potential to get stuck in a local minimum. This problem can be alleviated to some degree by including a random search component in the algorithm after a minimum has been achieved. These algorithms were found to be effective for suppressing combustion instabilities and will also be effective for optimization of combustion performance when the performance surface is not too rough.

The final type of algorithm considered was an LMS-based feedback configuration. The advantage of LMS is that it uses a stochastic gradient approach and is faster than the explicit gradient algorithms since there is no need to explicitly compute the gradient. The disadvantage is that the implicit gradient computation requires knowledge of a linearized model of the control-to-error response. This model must be determined in real time prior to starting the LMS iterations. This algorithm is well suited for the suppression of combustion

instabilities. The requirement of a linearized control-to-error model probably makes it unsuitable for combustion performance optimization, as the required model may be difficult to identify in real time.

Most control system analysis assumes that proportional actuation is available. In the combustion field, it is not uncommon to use on-off actuation in the form of fuel injectors. Control systems using on-off actuators have been effective in suppressing combustion instabilities, but there was no analysis of how such systems worked. We provide an analysis of the mechanism for achieving control using on-off actuators and validate the analysis with experimental results. Control systems using on-off actuation can be adaptively tuned using pattern search or explicit gradient algorithms with little modification. The filtered-E algorithm has also been effective when applied to on-off actuation, even though the implicit gradient is not quite correct.

The following chapters provide a detailed description and experimental results for each class of algorithm. Attention is paid to the type of a priori information required and how this information can be obtained. In addition, robustness issues are examined to understand potential problems that could occur in using these algorithms. Our experience in using these algorithms to suppress combustion instabilities has shown that all three types of algorithms are effective for adaptive control of these instabilities.

1.2 Graduate Students Supported:

Michael Vaudrey	PhD	2000
Xin Ming Huang	PhD	2001
Matthew Carson	M.S.	2001
Aaron Greenwood	M.S.	2002

1.3 Papers Published:

1. Vaudrey, M. A., W. T. Baumann, and W. R. Saunders, "Time-Averaged Gradient Control of Thermoacoustic Instabilities," accepted by *Journal of Propulsion and Power*.
2. Vaudrey, M.A., W.T. Baumann, and W.R. Saunders, "Stability and Operating Constraints of Adaptive Feedback Control" accepted by *Automatica*.
3. Vaudrey, M.A., W.T. Baumann, and W.R. Saunders, "Applying Adaptive LMS to Feedback Control of Thermoacoustic Instabilities" submitted to *IEEE Transactions on Control Systems Technology*, June, 2001.
4. Baumann, W.T, W.R. Saunders, and M.A. Carson, "Analysis and Experiments in Subharmonic Control of Thermoacoustic Instabilities," submitted to *ASME Journal of Dynamics, Measurement, and Control*, September, 2002.
5. C. A. Fannin, W. T. Baumann, and W. R. Saunders, "Thermoacoustic Stability Analysis for Multiport Fuel Injection in a Lean, Premixed Combustor," AIAA Paper 2000-0711, 38th Aerospace Sciences Meeting, Reno, NV, January 2000.

6. M. A. Vaudrey, W. R. Saunders, and W. T. Baumann, "An Investigation of Adaptive Signal Processing Approaches to Active Combustion Control," Proceedings of the Applied Vehicle Technology Panel Symposium on Active Control Technology for Enhanced Performance Operation Capabilities of Military Aircraft, Land Vehicles, and Sea Vehicles, Braunschweig, Germany, May, 2000.

2 Pattern Search Algorithms

2.1 Introduction

This chapter discusses the work done with algorithms that use a pattern search for the suppression of thermoacoustic instabilities. The algorithms were used to adapt the weights of a two-tap FIR filter. In general, these algorithms make perturbations along specified search directions, and compute the corresponding change in the cost function. The algorithms then make a decision about future perturbations and search directions based on these computations.

2.2 Theory

2.2.1 Hooke and Jeeves Pattern Search with Line Search

The first pattern search algorithm discussed here is that of Hooke and Jeeves. The original algorithm proposed by Hooke and Jeeves was a strict pattern search along the weight directions. The version incorporated in this work is an adapted version that uses both a pattern search and a line search along the successful directions. Bazaraa and Shetty[2] summarized the Hooke and Jeeves Pattern Search as follows:

Initialization Step: $\mathbf{d}_1, \dots, \mathbf{d}_n$ are the coordinate directions. The scalar $\varepsilon > 0$ is used to terminate the algorithm. Also choose an initial step size, $\Delta > \varepsilon$, and an acceleration factor, $\alpha > 0$. Choose a starting point \mathbf{x}_1 , let $\mathbf{y}_1 = \mathbf{x}_1$, let $k = j = 1$, and go to main step.

Main Step:

1. If $f(\mathbf{y}_j + \Delta \mathbf{d}_j) < f(\mathbf{y}_j)$, the trial is deemed a success; let $\mathbf{y}_{j+1} = \mathbf{y}_j + \Delta \mathbf{d}_j$, and go to step 2. If $f(\mathbf{y}_j + \Delta \mathbf{d}_j) \geq f(\mathbf{y}_j)$, the trial is deemed a failure. In this case, if $f(\mathbf{y}_j - \Delta \mathbf{d}_j) < f(\mathbf{y}_j)$, let $\mathbf{y}_{j+1} = \mathbf{y}_j - \Delta \mathbf{d}_j$, and go to step 2; if $f(\mathbf{y}_j - \Delta \mathbf{d}_j) \geq f(\mathbf{y}_j)$, let $\mathbf{y}_{j+1} = \mathbf{y}_j$, and go to step 2.
2. If $j < n$, replace j by $j + 1$, and repeat step 1. Otherwise go to step 3 if $f(\mathbf{y}_{n+1}) < f(\mathbf{x}_k)$, and go to step 4 if $f(\mathbf{y}_{n+1}) \geq f(\mathbf{x}_k)$.
3. Let $\mathbf{x}_{k+1} = \mathbf{y}_{n+1}$, and let $\mathbf{y}_1 = \mathbf{x}_{k+1} + \alpha(\mathbf{x}_{k+1} - \mathbf{x}_k)$. Replace k with $k + 1$, let $j = 1$, and go to step 1.
4. If $\Delta < \varepsilon$, stop; \mathbf{x}_k is the solution. Otherwise replace Δ by $\Delta/2$. Let $\mathbf{y}_1 = \mathbf{x}_k$, $\mathbf{x}_{k+1} = \mathbf{x}_k$, replace k by $k + 1$, let $j = 1$, and repeat step 1.

This algorithm can be thought of in two separate phases: an exploratory search (represented by steps 1 and 2 above) and a pattern search (step 3). The exploratory search successively perturbs along each weight direction and tests the resulting performance. The pattern search steps along the $(x_{k+1} - x_k)$ direction, or the direction between the last two points selected by the exploratory search. When perturbations along both the positive and negative weight direction do not result in enhanced performance, the perturbation size is decreased. When the perturbation size is less than a predetermined termination factor, ε , the algorithm stops and assumes an optimal solution.

2.2.2 Method of Rosenbrock

The method of Rosenbrock uses a pattern search technique to evaluate functional values along the search directions. An acceleration term is included to increase or decrease the step size as the algorithm progresses. Bazaraa and Shetty[2] summarized the Hooke and Jeeves Pattern Search as follows:

Initialization Step: $\mathbf{d}_1, \dots, \mathbf{d}_n$ are the coordinate directions and $\bar{\Delta}_1, \dots, \bar{\Delta}_n > 0$ be the initial step sizes along these directions. The scalar $\varepsilon > 0$ is used to terminate the algorithm. Let $\alpha > 1$ be the expansion factor and $\beta \in (-1, 0)$ be the contraction factor. Choose a starting point \mathbf{x}_1 , let $\mathbf{y}_1 = \mathbf{x}_1$, let $k = j = 1$, let $\Delta_j = \bar{\Delta}_j$ for each j , and go to main step.

Main Step:

1. If $f(y_j + \Delta_j d_j) < f(y_j)$, the trial is deemed a success; let $y_{j+1} = y_j + \Delta_j d_j$ and $\Delta_j = \alpha \Delta_j$. If $f(y_j + \Delta_j d_j) \geq f(y_j)$, the trial is deemed a failure; let $y_{j+1} = y_j$ and $\Delta_j = \beta \Delta_j$. If $j < n$, replace j by $j + 1$, and repeat step 1. Otherwise, if $j = n$, go to step 2.
2. If $f(y_{n+1}) < f(y_1)$, that is if any of the n trials of step 1 were successful, let $y_1 = y_{n+1}$, set $j = 1$, and repeat step 1. Consider the case when $f(y_{n+1}) = f(y_1)$, that is when each of the n trials in step 1 was a failure. If $f(y_{n+1}) < f(x_k)$, that is if at least one trial was successful during iteration k , go to step 3. If $f(y_{n+1}) = f(x_k)$, stop with x_k as an estimate for the optimal solution if $|\Delta_j| \leq \varepsilon$ for each j ; otherwise, let $y_1 = y_{n+1}$, $j = 1$, and go to step 1.
3. Let $x_{k+1} = y_{n+1}$. If $\|x_{k+1} - x_k\| < \varepsilon$, stop with x_{k+1} as an optimal solution.

Otherwise, compute $\lambda_1, \dots, \lambda_n$ from the relationship $x_{k+1} - x_k = \sum_{j=1}^n \lambda_j d_j$, form a new set of directions for the Gram-Schmidt procedure, let $\Delta_j = \bar{\Delta}_j$ for each j , let $y_1 = x_{k+1}$, replace k with $k + 1$, let $j = 1$, and repeat step 1.

The Gram-Schmidt procedure used in the Rosenbrock algorithm for calculating new directions can be expressed as:

$$a_j = \begin{cases} d_j & \text{if } \lambda_j = 0 \\ \sum_{i=j}^n \lambda_i d_i & \text{if } \lambda_j \neq 0 \end{cases} \quad (2.1)$$

$$b_j = \begin{cases} a_j & \text{if } j = 1 \\ a_j - \sum_{i=1}^{j-1} (a_j' \bar{d}_i) \bar{d}_i & \text{if } j \geq 2 \end{cases} \quad (2.2)$$

$$\bar{d}_j = \frac{b_j}{\|b_j\|} \quad (2.3)$$

This algorithm begins with searches along the weight directions. If a perturbation results in a lower MSE, the expansion term, α , is used to increase the step size in that direction. If a higher MSE results, the contraction term, β , reverses the sign of the perturbation and decreases the magnitude. This is repeated until a failure occurs along all directions, which leads to the development of new directions through the Gram-Schmidt procedure.

The Gram-Schmidt procedure takes the mutually orthogonal, linearly independent weight directions, d_1, \dots, d_n , and forms new directions $\bar{d}_1, \dots, \bar{d}_n$. The result is a new set of linearly independent orthogonal search directions. The Gram-Schmidt procedure forms a special case of conjugate directions, and thus the Rosenbrock algorithm behaves similar to a conjugate direction method in its convergence upon a quadratic function.

2.3 Robustness and Implementation Issues

For the above algorithms, the function that was to be minimized was the mean squared error, defined as

$$MSE = \frac{1}{N} \sum_{k=1}^N e^2(k) \quad (2.4)$$

where e is the error signal of interest and N is the number of samples over which the error is averaged.

For the pattern search algorithms, the MSE was made to be mean-zero data by subtracting the mean value from each element prior to squaring the signal. This was done because it was determined that a slight shift in the DC value could mask changes in the amplitude of the signal.

The Hooke and Jeeves algorithm was terminated if the MSE was below a threshold value. This was also incorporated to prevent the algorithm from making decisions based on signals not related to the initial limit cycle, such as noise from the pressure transducer.

The pattern search algorithms described above are functions of two very important parameters: the integration length, N , and the perturbation size, δ . As mentioned earlier, the integration length is the number of samples over which the MSE is averaged. The perturbation size corresponds to the change in the weights during the pattern search. These parameters are system-dependent, thus should be chosen automatically to insure the robustness of the minimization process.

The value of the MSE approaches the true mean value as the integration length approaches infinity. Because the algorithms are run in real-time and there is a desire to minimize convergence time, the goal is to find a minimum integration length while maintaining an accurate estimate of the MSE. This procedure involved reading the error signal and computing the MSE at each sample. The past n MSE values were then evaluated, where n is a specified evaluation range. When the maximum and minimum MSE values within this range fell within a certain percentage of each other, the current number of samples was considered an adequate integration length. If this condition was not met, the integration length was increased and the next MSE value was computed. The result of this is a transient response similar to a second order system. An example of this procedure is shown in Figure 2.1, which is a plot of the MSE calculation versus the sample number for an electronic simulation. At approximately 300 samples, the calculation of the MSE has reached a near steady-state value.

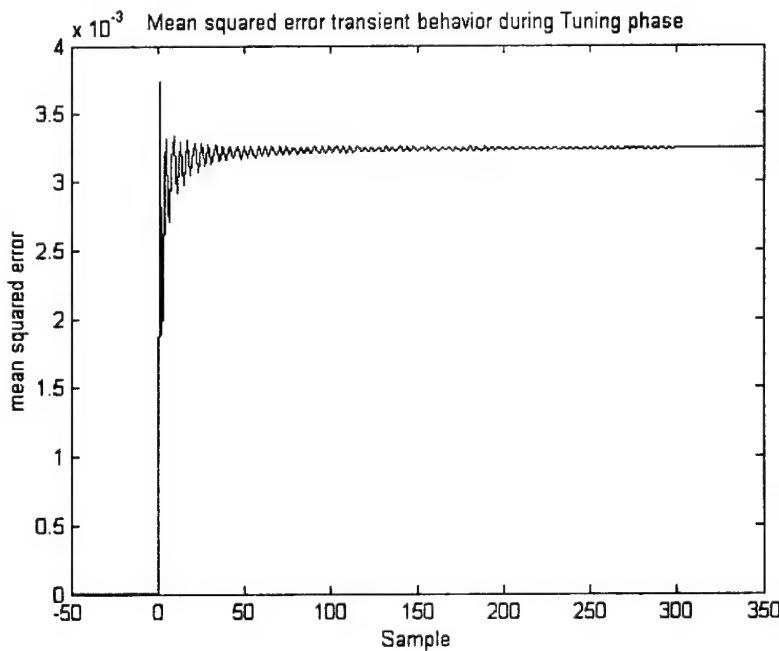


Figure 2.1 Mean squared error transient behavior during Tuning phase

It was necessary for the perturbation size to be large enough to cause a measurable change in the signal, yet not so large as to drive the system unstable once control was achieved. The initial perturbation size was chosen to be relatively small. The first weight was perturbed and the change in the MSE was measured. If the percent change was below a minimum value, the perturbation size was increased. If the change exceeded a maximum

value, it was decreased. Thus, when the change fell into a certain range, it was deemed acceptable. In practice, the acceptable range was a 20 – 25% change in the MSE.

2.4 Experimental Results

2.4.1 Experimental Setup

The algorithms discussed above were used in the stabilization of an open-closed tube combustor, commonly referred to as a Rijke tube. A ceramic honeycomb was used to stabilize a premixed air-methane flame located at one-half the length of the tube. The interaction between the heat release rate and the pressure fluctuations in the tube form a self-excited loop. Rayleigh's criterion predicts an instability of the second acoustic mode of the tube at a frequency of approximately 180 Hz, and that instability is observed experimentally.

A SenSym resistive-based pressure transducer was used to collect sound pressure data from this 180 Hz instability. This signal was sent through a strain gauge amplifier, then was band-pass filtered between 160 and 200 Hz. The signal was then sent to a dSpace DS1103 DSP board, sampling at 1600 Hz, and used by the minimization algorithms to determine adaptive filter coefficients. The output was passed through a smoothing filter at 185 Hz and into an amplifier. The amplified signal was delivered to a 3" speaker that was used to control the instability in the tube. A schematic of this system can be seen in Figure 2.2.

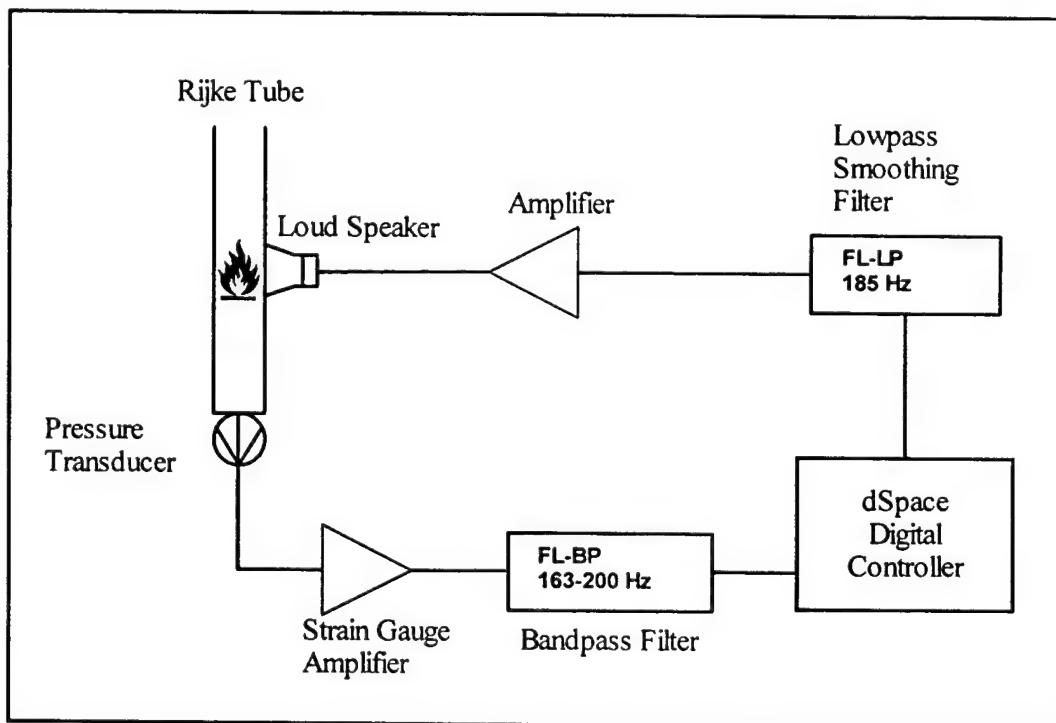


Figure 2.2 Rijke Tube Experimental Set-Up

Control tests were done in one of two actuation methods: proportional, which will be discussed later in the chapter, or pulsed (on-off) actuation signals, which is discussed in

more detail in Chapter 5. Proportional actuation involved a control signal that is proportional to the input pressure signal. Thus, the control effort is initially high, and decreases as control is achieved. Pulsed control, on the other hand, involves control signals of fixed amplitude. The pulsed experiments also include the use of sub-harmonic control. This is done by forcing the actuator at frequencies that are integer divisions of the fundamental instability frequency. Control is still possible because these signals contain harmonic component at the instability frequency. Sub-harmonic control may be used to extend component life by exposing the actuator to fewer cycles.

2.4.2 Performance using Proportional Actuation on Rijke tube

There were three separate operating conditions on the Rijke tube that were explored, each have a total flow of approximately 128 cc/sec. The first was a condition in which the tube was at the lower end of the instability range. The equivalence ratio for this case was, $\Phi = 0.544$. The second condition was at a higher equivalence ratio of $\Phi = 0.582$. The final condition was at an equivalence ratio, $\Phi = 0.641$, could not be successfully control due to restrictions in the actuator authority which are discussed at the end of this section. Because of this authority issue, higher equivalence ratio conditions were not investigated.

Because the Rijke tube is basically a single frequency noise cancellation problem, only two filter coefficients are required to control the phase and magnitude of the filter. This is shown through the equations for the magnitude and phase of the filter.

$$\angle G = \tan^{-1} \left(\frac{[-w_0 \sin(\omega T) + w_1 \sin(2\omega T)]}{[w_0 \cos(\omega T) + w_1 \cos(2\omega T)]} \right) \quad (2.5)$$

$$|G| = \sqrt{[w_0 \cos(\omega T) + w_1 \cos(2\omega T)]^2 + [w_0 \sin(\omega T) + w_1 \sin(2\omega T)]^2} \quad (2.6)$$

It is desired to constrain the magnitude and phase of the filter at the instability frequency such that the control signal will destructively interfere with the pressure oscillations generated from the instability. Because there is a single frequency that is relevant in this case, there is only a single magnitude and phase to constrain. As can be seen from the above equations, this results in the need for a filter with only two taps. Thus, unless otherwise stated, it can be assumed that two weights have been used to obtain the following results.

2.4.2.1 Low Equivalence Ratio Case

The first of the operating conditions for the Rijke tube was a relatively low equivalence ratio of 0.544 with limit cycle frequency of 178 Hz at a level of 4.65 dBVRms. The power spectrum of this operating condition can be seen in Figure 2.3 and traces of the pressure oscillations are shown in Figure 2.4.

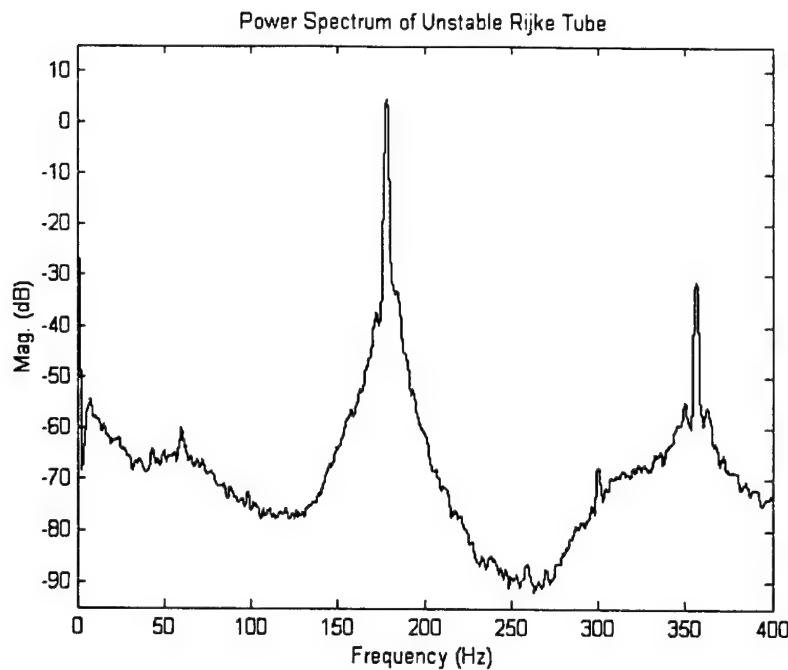


Figure 2.3 Power spectrum of Rijke tube, $\Phi = 0.544$

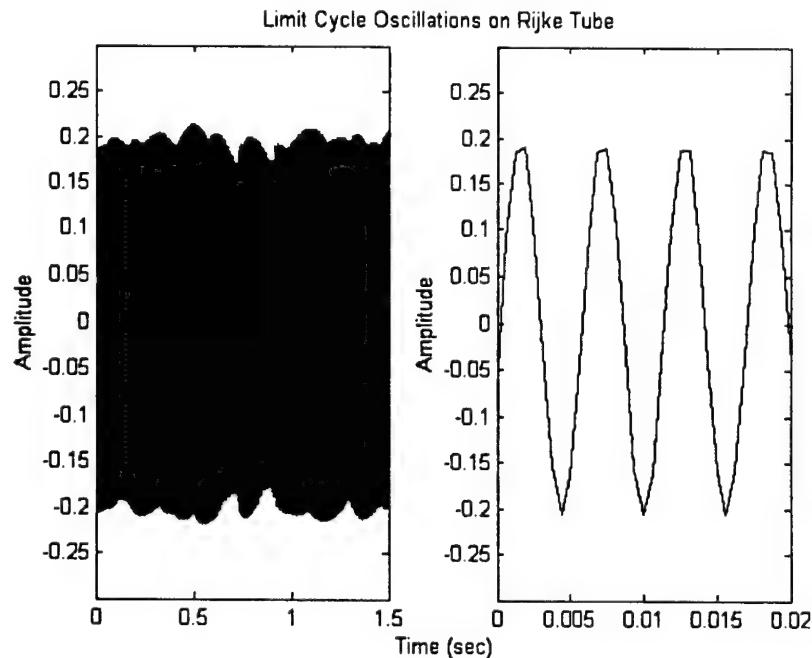


Figure 2.4 Limit cycle oscillations of Rijke Tube, $\Phi = 0.544$

The tuning phase that was described in Section 2.3 was used to determine acceptable values of δ , and N , the results of which can be seen in Table 2-1.

Table 2-1: Parameter Values for the $\Phi = 0.544$

Equivalence Ratio, Φ	0.544
Perturbation Size, δ	0.1644
Integration Length, N	402 samples

For each algorithm, data for the decay envelope, the magnitude and phase of the filter, and the power spectra of the controlled pressure in the tube were collected.

The results for the first of the pattern search algorithms, Hooke and Jeeves, can be seen in Figure 2.5 and Figure 2.6.

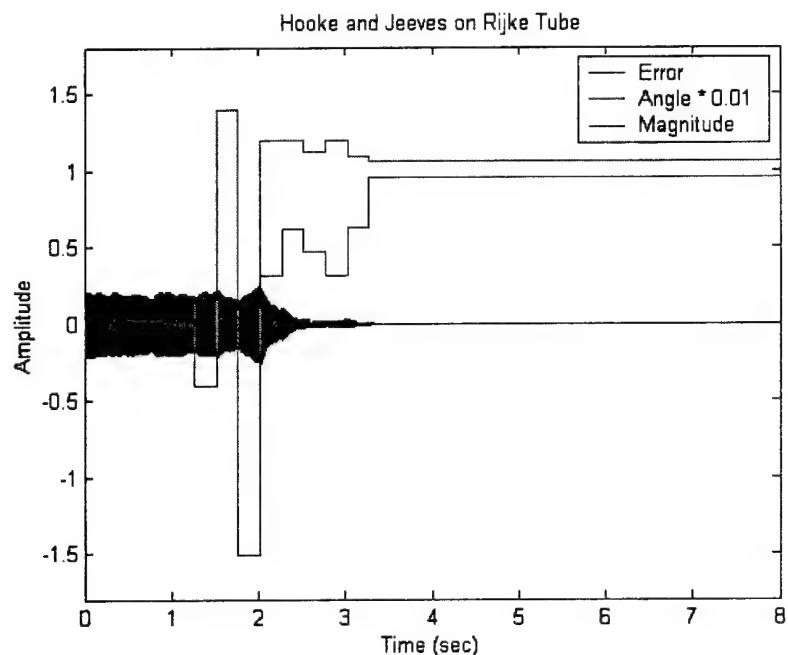


Figure 2.5 Phase and magnitude of filter for Hooke and Jeeves algorithm, $\Phi = 0.544$

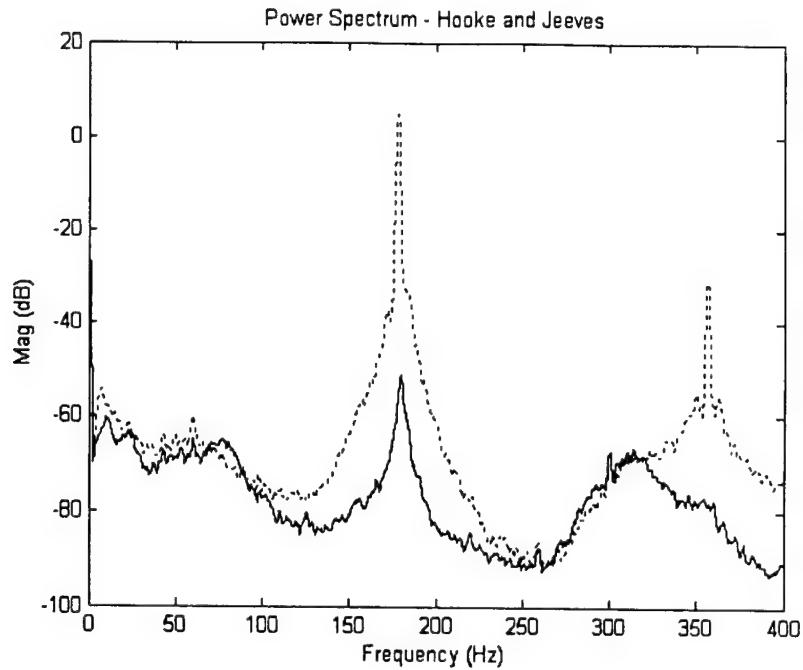


Figure 2.6 Power spectrum of Hooke and Jeeves-controlled Rijke Tube, $\Phi = 0.544$

The transient response for the Hooke and Jeeves pattern search does not show a smooth convergence. This is because the algorithm is systematically searching the weight space, and thus may take steps in the non-optimal direction.

The results for the final pattern search algorithm, Rosenbrock, are shown in Figure 2.7 and Figure 2.8.

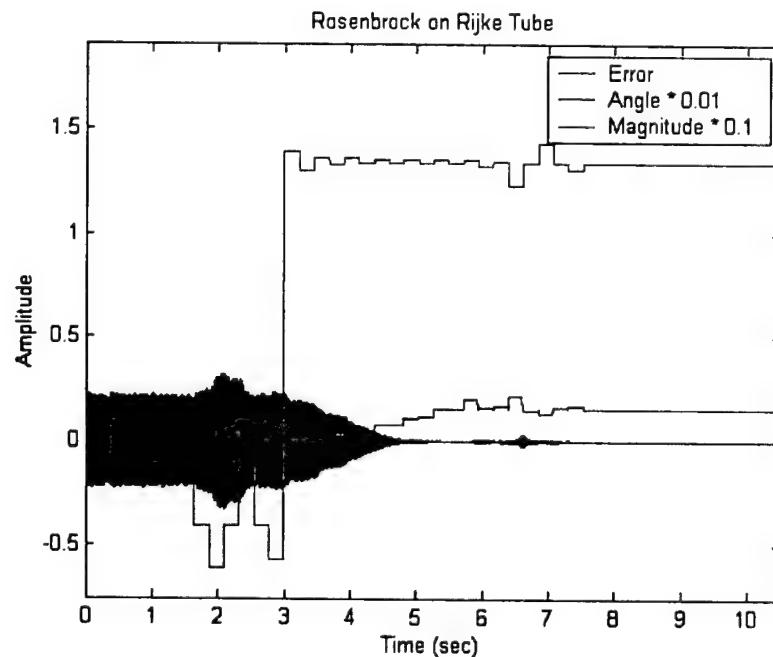


Figure 2.7 Phase and magnitude of filter for Rosenbrock algorithm, $\Phi = 0.544$

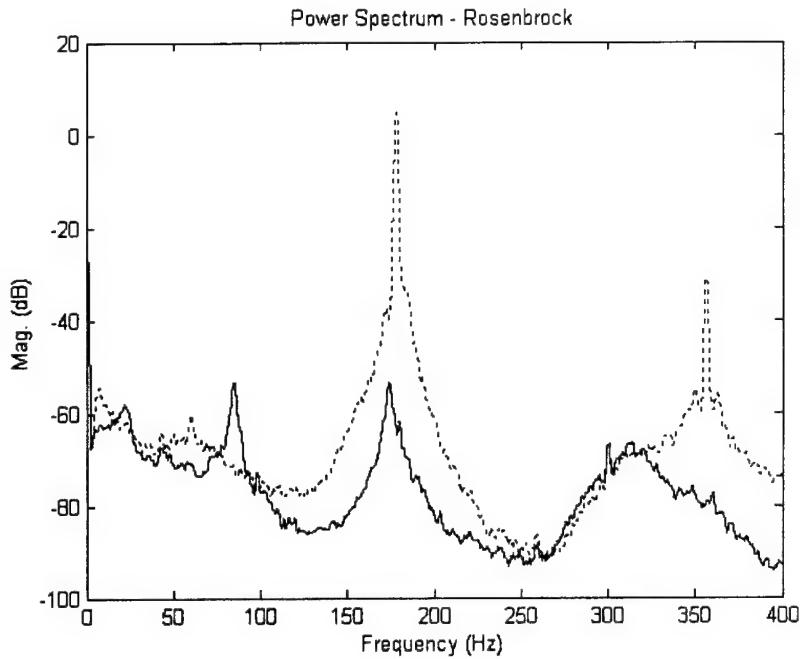


Figure 2.8 Power spectrum of Rosenbrock-controlled Rijke Tube, $\Phi = 0.544$

Because both are pattern searches, the statements related to the transient response of the Hooke and Jeeves algorithm also apply to the Rosenbrock algorithm. The convergence times, ultimate attenuation, and the phase and magnitude of the controllers for each of the algorithms at this operating condition are tabulated in Table 2-2. For the low equivalence ratio case, each algorithm was successful in gaining control of the thermoacoustic instability. On average, there was about 57 dB of attenuation and a convergence time of approximately 3 seconds.

Table 2-2: Convergence times and ultimate attenuation at limit cycle frequency for $\Phi = 0.544$

Algorithm	Convergence Time	Ultimate Attenuation	Filter Magnitude	Filter Phase
Hooke and Jeeves	2.219 sec.	56.1 dB	0.95	106°
Rosenbrock	3.852 sec.	58.2 dB	1.55	134°

2.4.2.2 Medium Equivalence Ratio Case

The second condition was at an equivalence ratio of 0.582. Figure 2.9 shows the power spectrum of this condition with the instability at nearly 180 Hz and a magnitude of approximately 8 dBVRms, while Figure 2.10 shows the envelope of the oscillations.

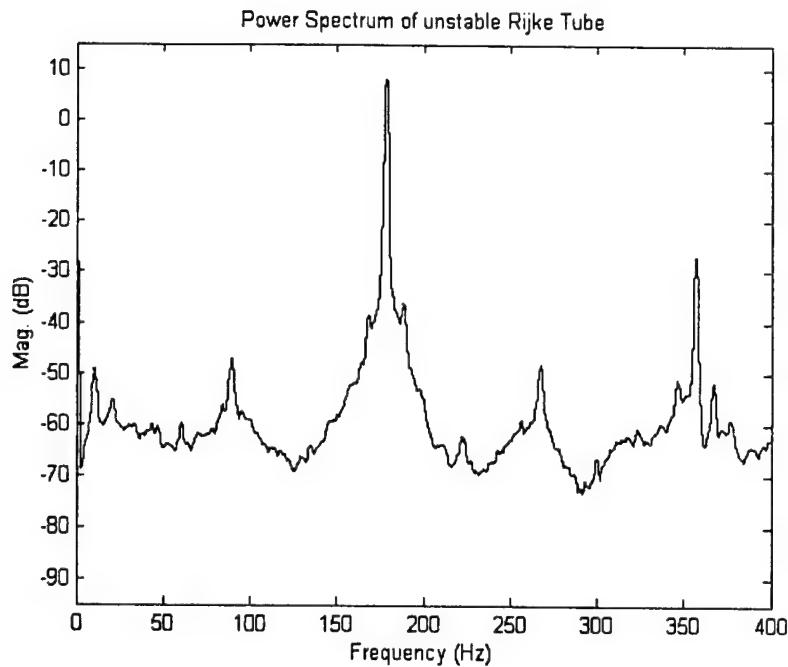


Figure 2.9 Power spectrum of Rijke Tube, $\Phi = 0.582$

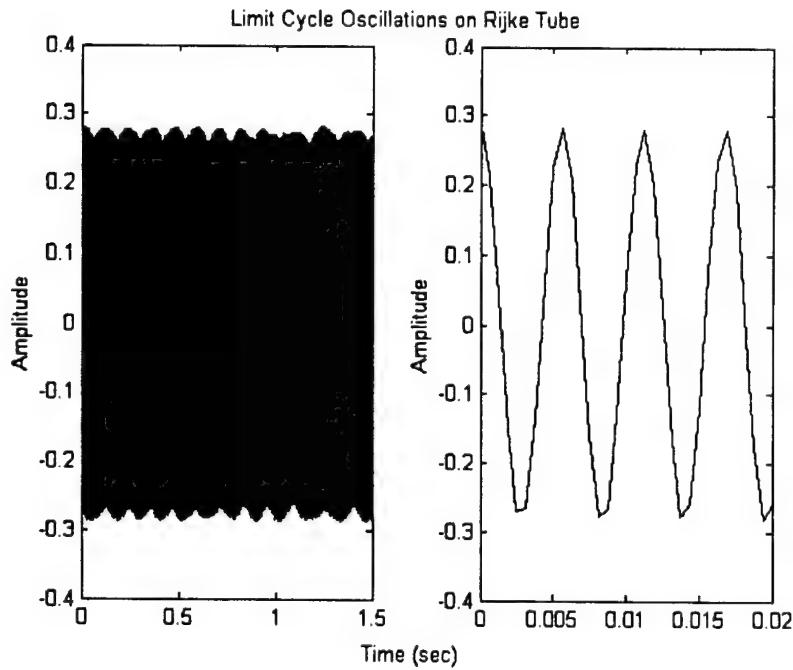


Figure 2.10 Limit cycle oscillations of Rijke Tube, $\Phi = 0.582$

The tuning phase was run for this operating condition and the results are given in Table 2-3 and the resulting convergence times for this condition are given in Table 2-4.

Table 2-3: Parameter Values for the $\Phi = 0.582$

Equivalence Ratio, Φ	0.582
Perturbation Size, δ	0.1621
Integration Length, N	361 samples

Raising the equivalence ratio to 0.582 did not negatively affect the performance of the algorithms. The results for this case, in fact, are very similar to that of the previous case. The results are shown in Figure 2.11 through Figure 2.14, encompassing both algorithms.

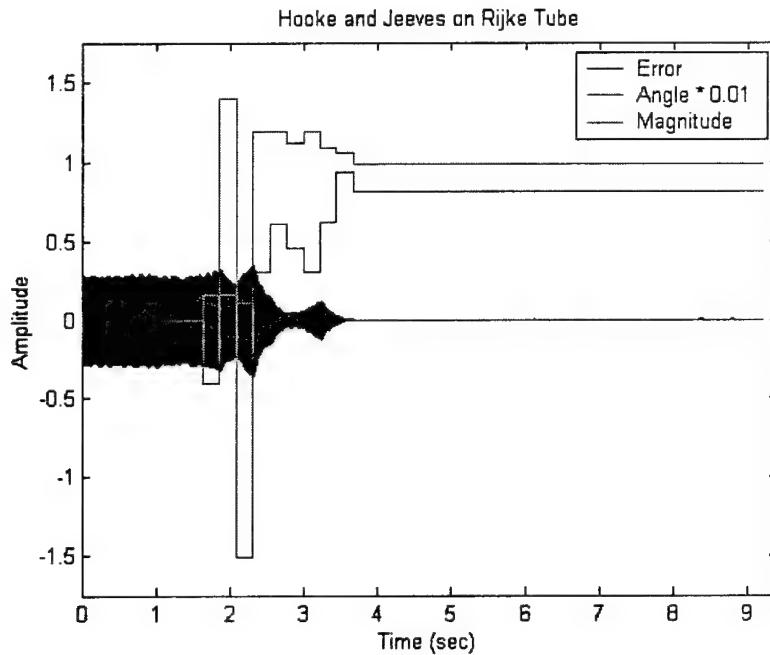


Figure 2.11 Phase and magnitude of filter Hooke and Jeeves algorithm, $\Phi = 0.582$

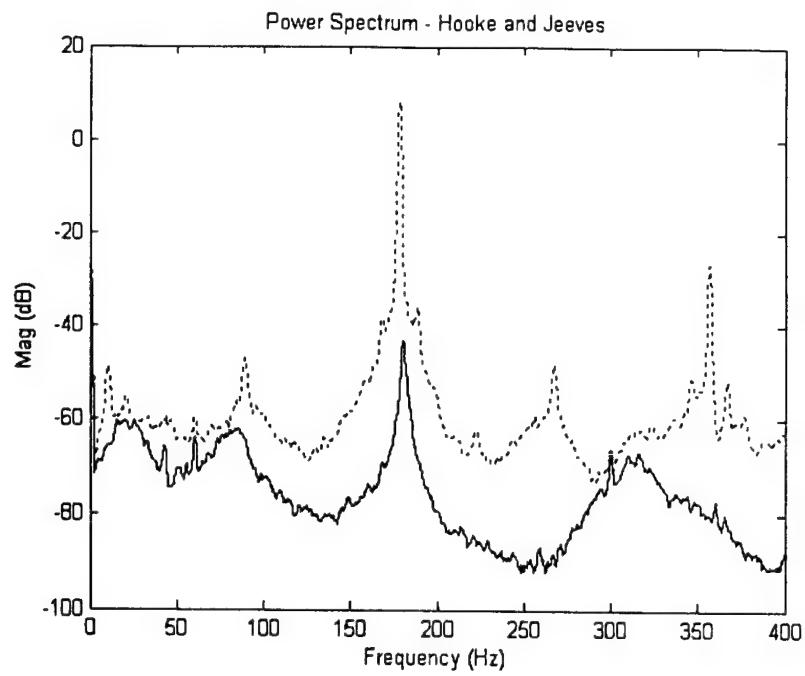


Figure 2.12 Power spectrum of Hooke and Jeeves-controlled Rijke Tube, $\Phi = 0.582$

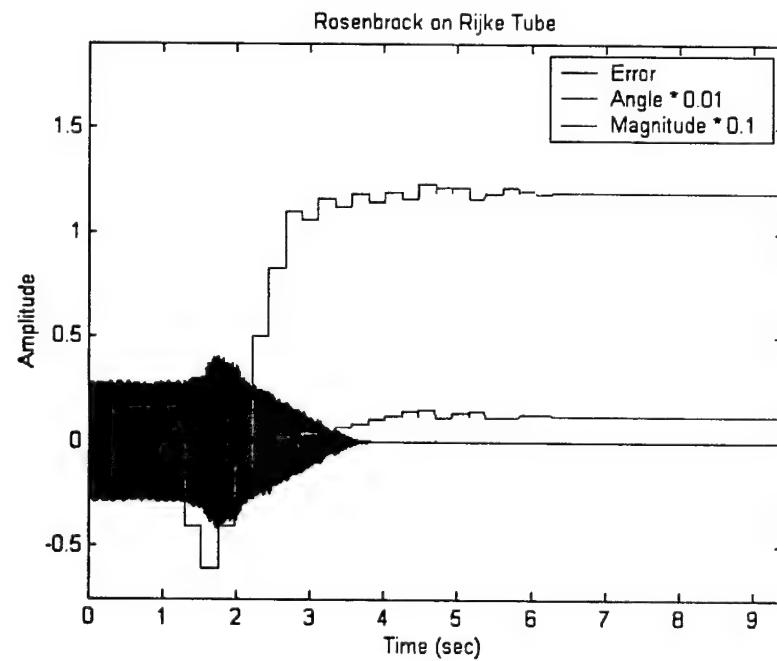


Figure 2.13 Phase and magnitude of filter Rosenbrock algorithm, $\Phi = 0.582$

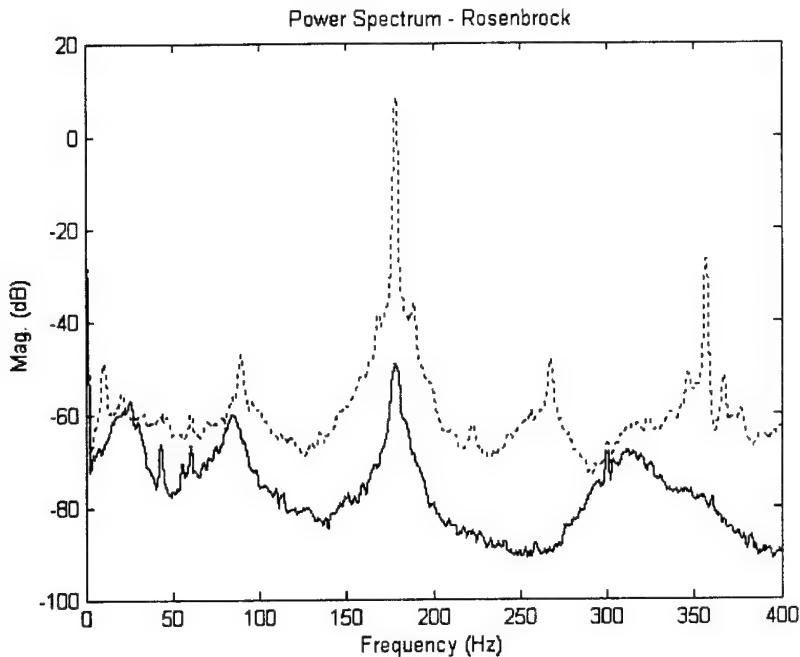


Figure 2.14 Power spectrum of Rosenbrock-controlled Rijke Tube, $\Phi = 0.582$

The results shown in the above figures are tabulated in Table 2-4, with the convergence time, ultimate attenuation, and magnitude and phase of the controller.

Table 2-4: Convergence times and ultimate attenuation at limit cycle frequency for the $\Phi = 0.582$

Algorithm	Convergence Time	Ultimate Attenuation	Magnitude	Phase
Hooke and Jeeves	2.075 sec.	50.8 dB	0.81	99°
Rosenbrock	2.726 sec.	57.1 dB	1.25	119°

The magnitudes and phases that were developed by the algorithms for this operating condition are similar to those for the previous condition, and the attenuation was approximately 55 dB. The average convergence time for this case was 2.4 seconds, which is slightly lower than the low equivalence ratio case. This is most likely due to a shorter integration length for this condition.

2.4.2.3 High Equivalence Ratio Case

The final condition that was tested on the Rijke tube was a high equivalence ratio of 0.641. This equivalence ratio produced the power spectrum in Figure 2.15 and the oscillation trace in Figure 2.16.

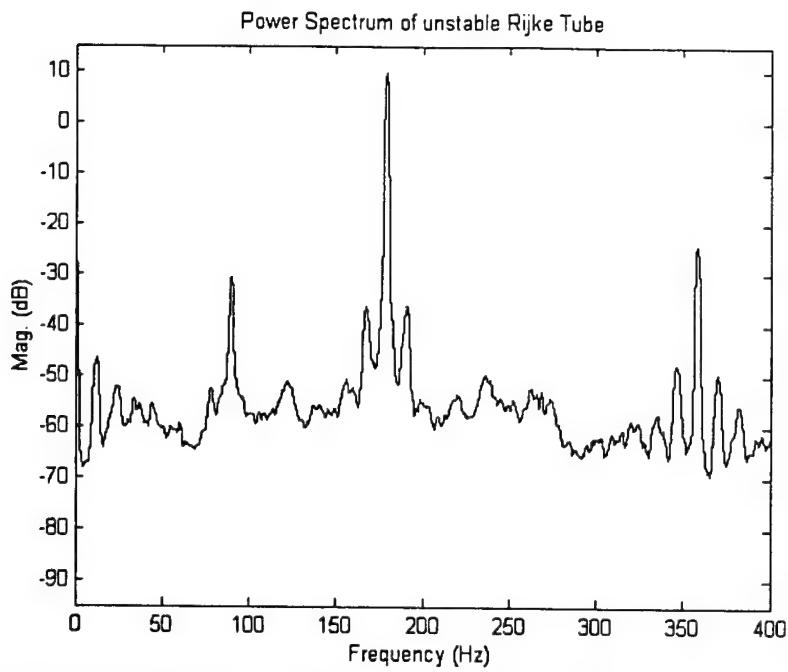


Figure 2.15 Power spectrum of Rijke Tube, $\Phi = 0.641$

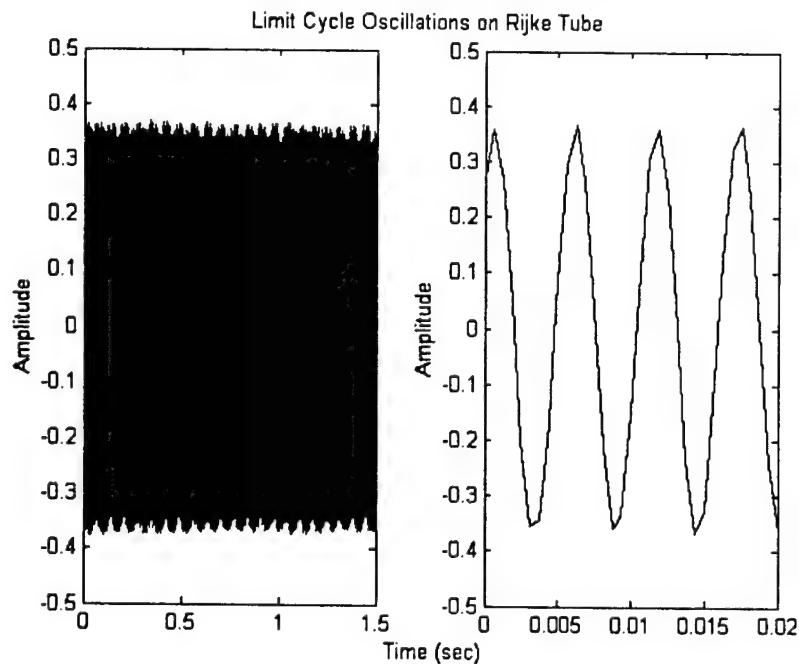


Figure 2.16 Limit cycle oscillations of Rijke Tube, $\Phi = 0.641$

The tuning phase was once again run and Table 2-5 outlines the results. Table 2-6 presents the convergence times of each algorithm for the high equivalence ratio operating condition.

Table 2-5: Parameter Values for the $\Phi = 0.641$

Equivalence Ratio, Φ	0.641
Perturbation Size, δ	0.2736
Integration Length, N	341 samples

Results for the Hooke and Jeeves algorithm on the high equivalence ratio condition are presented in Figure 2.17 and Figure 2.18. Note that control cannot be maintained for this condition. Initially, control is achieved, but then the system produces a series of bursts in the pressure signal. This phenomenon has been documented previously in Vaudrey [1], which states that the bursting is due to a combination of lack of actuator authority and a changing plant. As the flame becomes stable the heat release increases, which drives the poles further into the right half of the s-plane. As the leaner conditions, there is still enough authority to maintain control over this increasingly unstable plant, but at richer conditions the actuator lacks this authority. As the system begins to lose control and become unstable, the heat release decreases. This allows the actuator to regain control, and the process repeats itself in the form of a series of bursts in the combustor pressure.

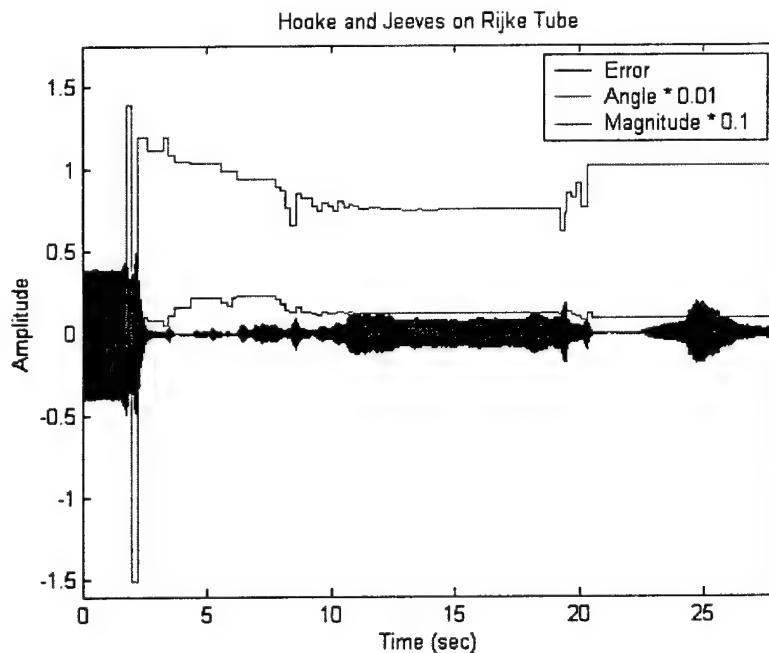


Figure 2.17 Phase and magnitude of filter Hooke and Jeeves algorithm, $\Phi = 0.641$

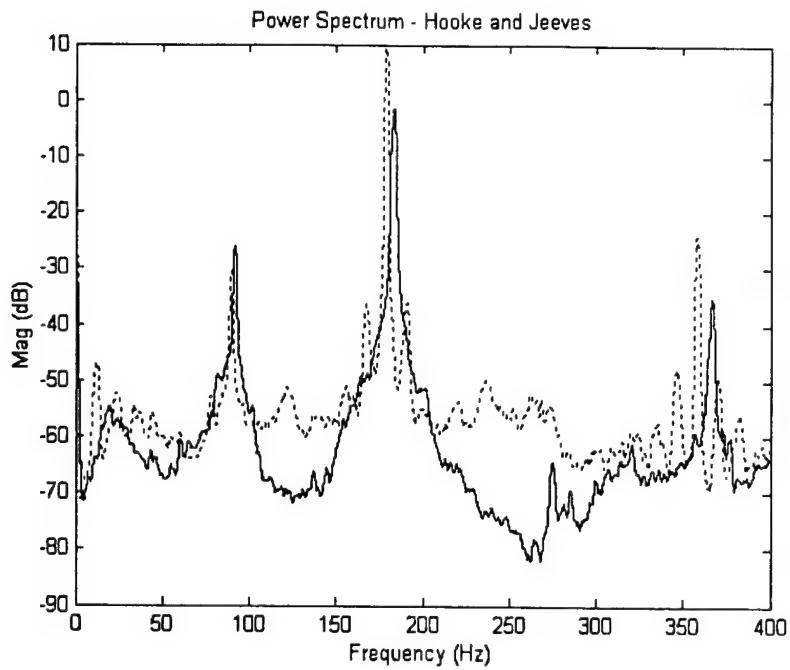


Figure 2.18 Power spectrum of Hooke and Jeeves-controlled Rijke Tube, $\Phi = 0.641$

Rosenbrock results are shown in Figure 2.19 and Figure 2.20.

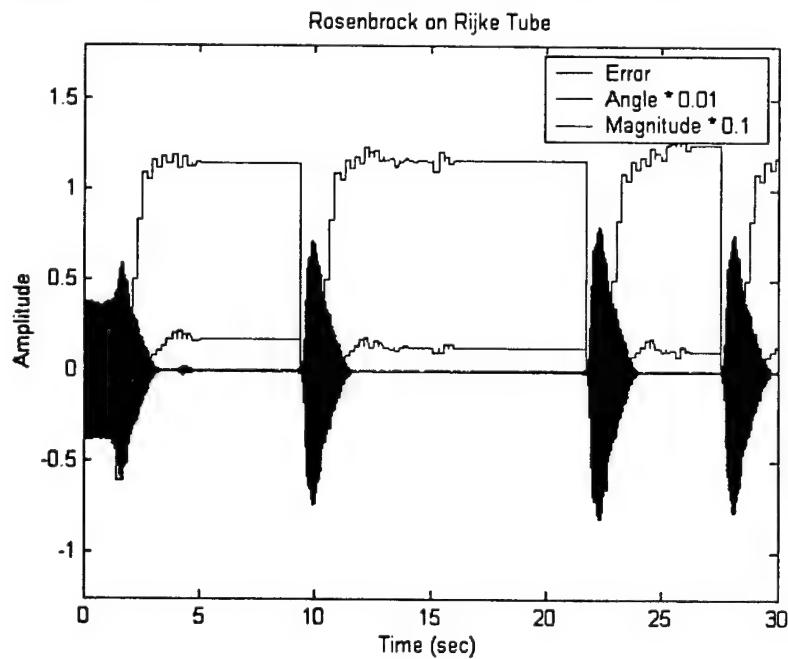


Figure 2.19 Phase and magnitude of filter Rosenbrock algorithm, $\Phi = 0.641$

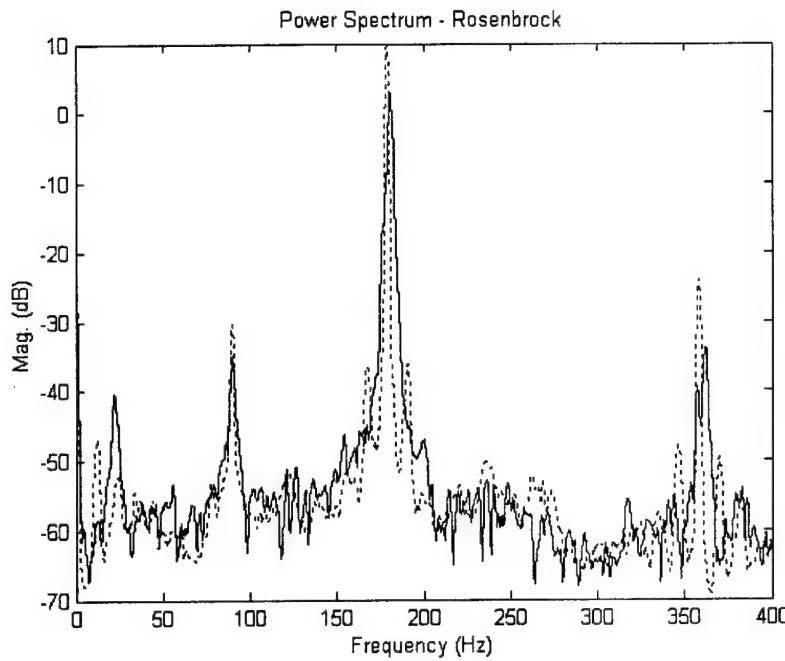


Figure 2.20 Power spectrum of Rosenbrock-controlled Rijke Tube, $\Phi = 0.641$

The transient response for the Rosenbrock algorithm shows extreme bursts in the pressure signal. The Rosenbrock algorithm searches the weight space and settles into a region that produces the lowest MSE of the pressure signal. The algorithm will begin to search again if a measurable increase in the MSE occurs. This is the reason that the bursts for this algorithm are so severe.

Table 2-6: Convergence times and ultimate attenuation at limit cycle frequency for the $\Phi = 0.641$

Algorithm	Convergence Time	Ultimate Attenuation
Hooke and Jeeves	2.193 sec.	11.0 dB
Rosenbrock	2.379 sec.	6.9 dB

2.4.2.4 Proportional Actuation Discussion

At this point, there should be a short discussion on the results of the tuning for each of the above conditions. It can be seen that the integration length, N , decreases with increasing equivalence ratio, from 402 for the low case to 341 for the high case. This is the expected result considering the plots of the limit cycling pressure signal. Figure 2.4 shows the trace of the low equivalence ratio case, and it shows that the amplitude of the oscillations is far from uniform. This is compared to Figure 2.10 for the medium case and Figure 2.16 for the high case, in which the envelope is much less sporadic. Thus, it can be concluded that the low equivalence ratio case should require the longest averaging time.

One would expect that the perturbation size, δ , would increase as the equivalence ratio increased. This is due to the fact that with an increase in equivalence ratio comes an increase in the limit cycle amplitude, thus a larger perturbation is required to achieve a measurable change in the MSE. The perturbation size stayed relatively constant from the low to medium case, but did increase substantially for the high case.

The Hooke and Jeeves algorithm has outperformed the Rosenbrock algorithm with respect to convergence time in the Rijke tube tests. The Hooke and Jeeves algorithm also was better suited for situations in which actuator authority was an issue, such as the $\Phi = 0.641$ case. When actuator authority was not an issue, the Rosenbrock algorithm did produce a slight advantage in the attenuation of the pressure signal.

This work has shown that the two pattern search algorithms that were investigated converged quickly and could maintain an acceptable level of attenuation.

2.5 References

- [1] M. A. Vaudrey. "Active Control Methods for Non-linear Self-excited Systems," PhD Dissertation, Virginia Tech, 2001.
- [2] M. S. Bazaraa, C. M. Shetty, *Nonlinear Programming: Theory and Algorithms*, John Wiley and Sons, Inc., New York, NY, 1979.
- [3] B. Widrow, S. D. Stearns, *Adaptive Signal Processing*, Prentice Hall, Inc., Englewood Cliffs, NJ, 1985.
- [4] D. M. Himmelblau, *Applied Nonlinear Programming*, McGraw-Hill, Inc., New York, NY, 1972.
- [5] K. McManus, T. Poinsot, S.M. Candel, "A Review of Active Control of Combustion Instabilities," *Prog. Energy Combust. Sci.* 19, 1993, pp. 1-29.
- [6] R.L. Raun, M.W. Beckstead, J.C. Finlinson, K.P. Brooks, "A Review of Rijke Tubes, Rijke Burners and Related Devices," *Prog. Energy Combust. Sci.* 19, 1993, pp. 313 – 364

3 Explicit Gradient Algorithms

3.1 Introduction

This chapter discusses the work done with explicit gradient algorithms. These algorithms are based on the method of steepest descent. In general, the explicit gradient algorithms perturb the filter weights, calculate the gradient of the cost function, and descend in the gradient direction using a specific weight update equation.

3.2 Theory

3.2.1 Time Averaged Gradient (TAG)

The TAG algorithm performs a gradient search by the method of steepest descent for which the basic weight update equation is:

$$w_n(k+1) = w_n(k) - \mu \xi'(w_n(k)) \quad (3.1)$$

where w_n is the n^{th} filter coefficient and the gradient, ξ' , is calculated from:

$$\xi'(w_n) = \frac{\xi(w_n) - \xi(w_n - \delta)}{\delta}. \quad (3.2)$$

The cost function, ξ , is the mean squared error (MSE) of the signal:

$$\xi(w_n) = \frac{1}{N} \sum_{k=0}^N e^2(k). \quad (3.3)$$

The convergence parameter, μ , is a system dependent value that controls the speed of adaptation and stability of the filter.

A more detailed description of this algorithm can be found in Widrow [3] and experimental results were presented by Vaudrey[1].

3.2.2 Gradient Descent with Line Search

This algorithm is an extension of the TAG algorithm, with the addition of a line search along the gradient direction. The gradient is calculated using the same procedure described above, but instead of taking a fixed step, μ , along the gradient, the algorithm searches along the gradient direction for the optimal step size. This is done by taking successively larger step sizes and evaluating the corresponding MSE. When an increase in step size no longer results in increased performance, a new gradient is calculated, and the procedure is repeated.

3.2.3 Fletcher and Reeves Conjugate Gradient

The conjugate gradient method of Fletcher and Reeves creates search directions that are a linear combination of the steepest decent direction and previous search directions. Weighting factors are applied such that the search directions are conjugate. These factors are ratios of the present and past squared norms of the gradient. Bazaraa and Shetty[2] summarized the Fletcher and Reeves Conjugate Gradient as follows:

Initialization Step: The scalar $\varepsilon > 0$ terminates the algorithm and the initial point is x_1 . Let $y_1 = x_1$, $d_1 = -\nabla f(y_1)$, $k = j = 1$, and go to the main step

Main Step:

1. If $\|\nabla f(y_j)\| < \varepsilon$, stop. Otherwise, let λ_j be the optimal solution to the problem to minimize $f(y_j + \lambda_j d_j)$ subject to $\lambda_j \geq 0$, and let $y_{j+1} = y_j + \lambda_j d_j$. If $j < n$, go to step 2, otherwise go to step 3.
2. Let $d_{j+1} = -\nabla f(y_{j+1}) + \alpha_j d_j$, where $\alpha_j = \frac{\|\nabla f(y_{j+1})\|^2}{\|\nabla f(y_j)\|^2}$. Replace j by $j + 1$, and go to step 1.
3. Let $y_1 = x_{k+1} = y_{n+1}$, and let $d_1 = -\nabla f(y_1)$. Let $j = 1$, replace k by $k + 1$, and go to step 1.

The concept of conjugacy is very important in unconstrained optimization problems. If the function to be minimized is quadratic, conjugate search directions guarantee that convergence will occur in at most n steps, where n is the number of parameters being adapted.

3.3 Robustness and Implementation Issues

The gradient descent algorithms used the same means of computing the MSE as mentioned in the previous chapter for the pattern search algorithms. These algorithms also used the same procedure for computing the integration length, N , and the perturbation size, δ . The TAG algorithm is also a function of the convergence parameter, μ , which was also calculated during the tuning phase.

Widrow [3] gives the bounds of μ for the LMS algorithm as

$$\frac{1}{\lambda_{\max}} > \mu > 0 \quad (3.4)$$

where λ_{\max} is the largest eigenvalue of the input correlation matrix, R. The R matrix is defined as

$$R = E[X_k X_k^T] \quad (3.5)$$

where X_k is the input matrix.

It can then be stated that λ_{\max} cannot be greater than the trace of R ($\text{tr}[R]$), which is the sum of the diagonal elements of R. From the definition of R,

$$\text{tr}[R] = (L + 1)E[x_k^2] \quad (3.6)$$

where L is the number of inputs. Thus the bounds of μ can be stated as

$$0 < \mu < \frac{1}{(L + 1)(\text{signal power})} \quad (3.7)$$

In practice, the signal power was calculated identical to the MSE mentioned above, and the convergence parameter used was 25% of the upper bound given above. Because this expression is derived from work on the LMS algorithm and assumes a transversal filter, this approach is valid when used to adapt an FIR filter. This approach was used for this work because these algorithms were indeed used to adapt an FIR filter, but care must be taken if a different control structure is being employed. These bounds may no longer be valid.

3.4 Experimental Results

3.4.1 Performance using Proportional Actuation on Rijke tube

The experimental set-up that was outlined in the previous chapter was again used for tests of the explicit gradient algorithms. There were, again, three separate operating conditions that were employed during the tests.

3.4.1.1 Low Equivalence Ratio Case

The transient response, phase and magnitude, and power spectrum for the TAG algorithm control at the $\Phi = 0.544$ equivalence ratio condition are shown in Figure 3.1 and Figure 3.2.

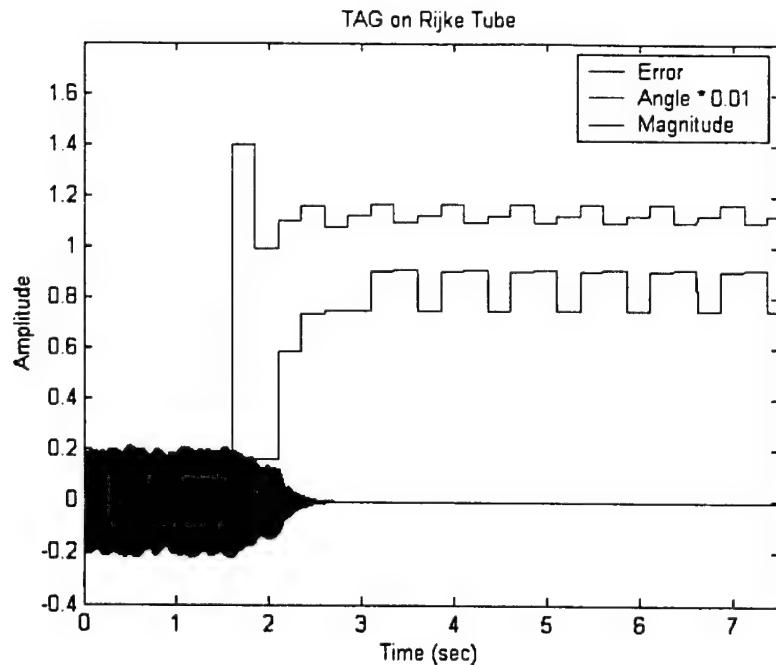


Figure 3.1 Phase and magnitude of filter for TAG algorithm, $\Phi = 0.544$

Refer to equations 2.5 and 2.6 for the conversion from filter weights to angle and phase of the filter.

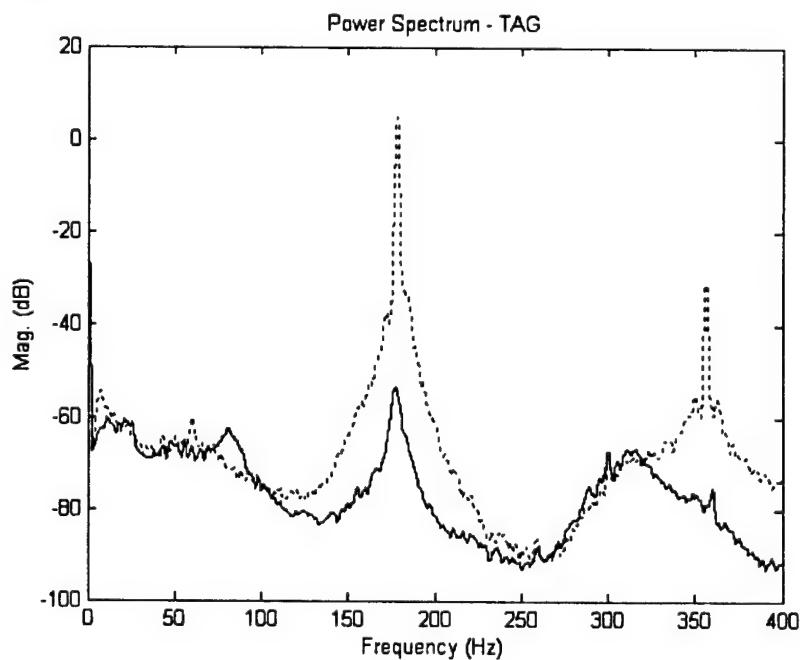


Figure 3.2 Power spectrum of TAG-controlled Rijke tube, $\Phi = 0.544$

The TAG algorithm was able to produce satisfactory attenuation of the pressure signal, along with a reasonable convergence time. Figure 3.1 shows that the magnitude and phase of the controller do not reach a steady-state value. This is because the algorithm

continuously perturbs the system to calculate a gradient, which causes continuous, small fluctuations in the magnitude and phase.

The results of the Gradient algorithm for this same operating condition are presented in Figure 3.3 and Figure 3.4.

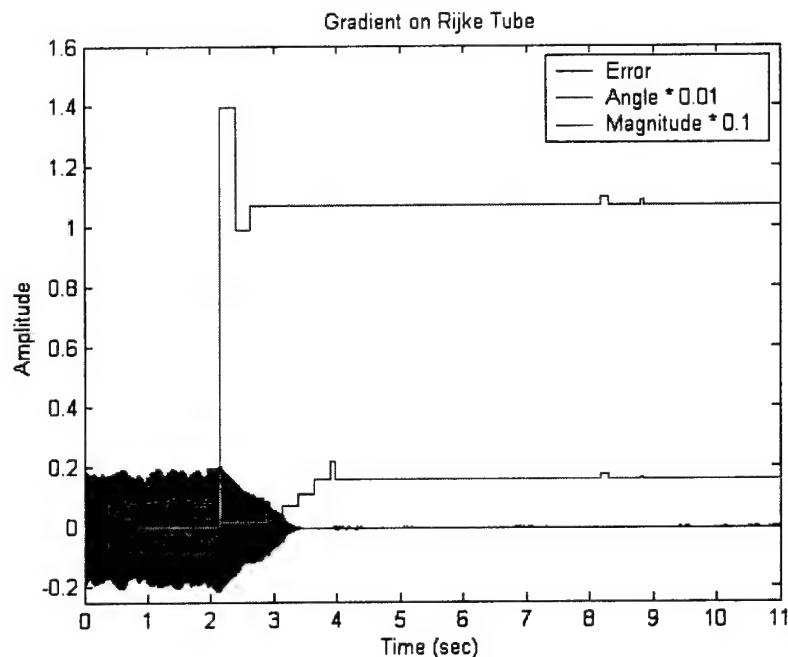


Figure 3.3 Phase and magnitude of filter for Gradient algorithm, $\Phi = 0.544$

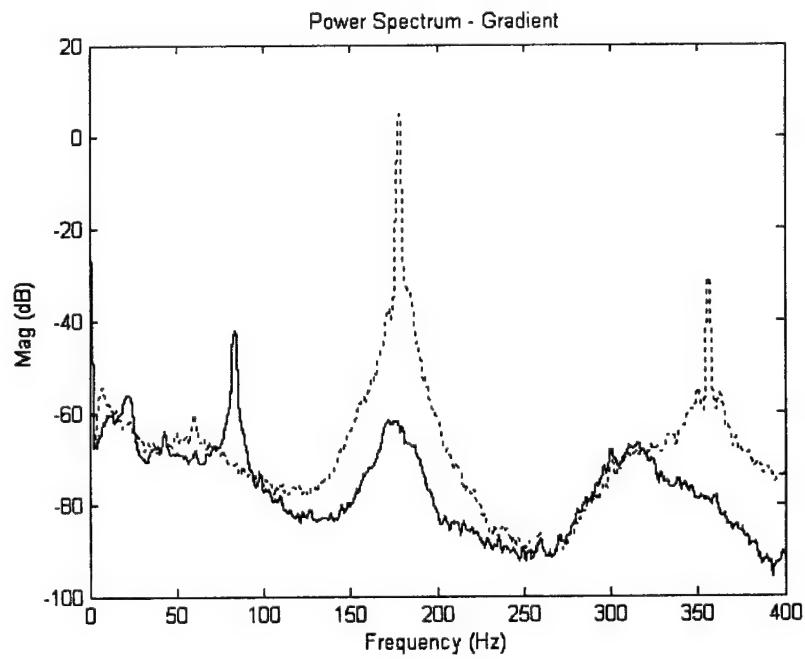


Figure 3.4 Power spectrum of Gradient-controlled Rijke Tube, $\Phi = 0.544$

The Gradient algorithm, too, performed well at this condition. It is important to note that the dominant energy component in the power spectrum has shifted from the initial limit cycling frequency of approximately 180 Hz to the second sub-harmonic frequency of 90 Hz. This is a result of using a bandpass filter on the pressure signal. A narrow bandpass was used, and as a result the information at the subharmonics is filtered. The Gradient algorithm does not see the growth at this frequency, thus does not adapt to reduce this signal component.

Results for the Conjugate algorithm are shown in Figure 3.5 and Figure 3.6.

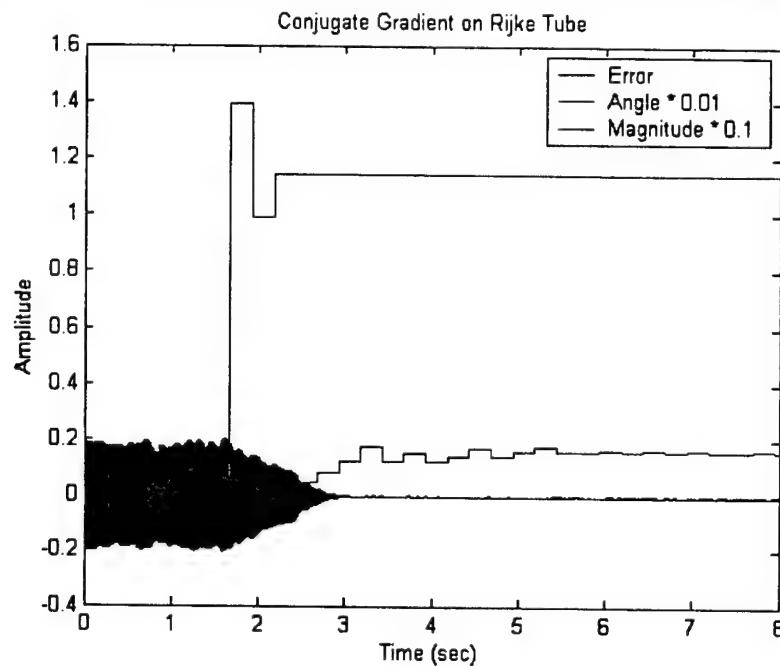


Figure 3.5 Phase and magnitude of filter for Conjugate algorithm, $\Phi = 0.544$

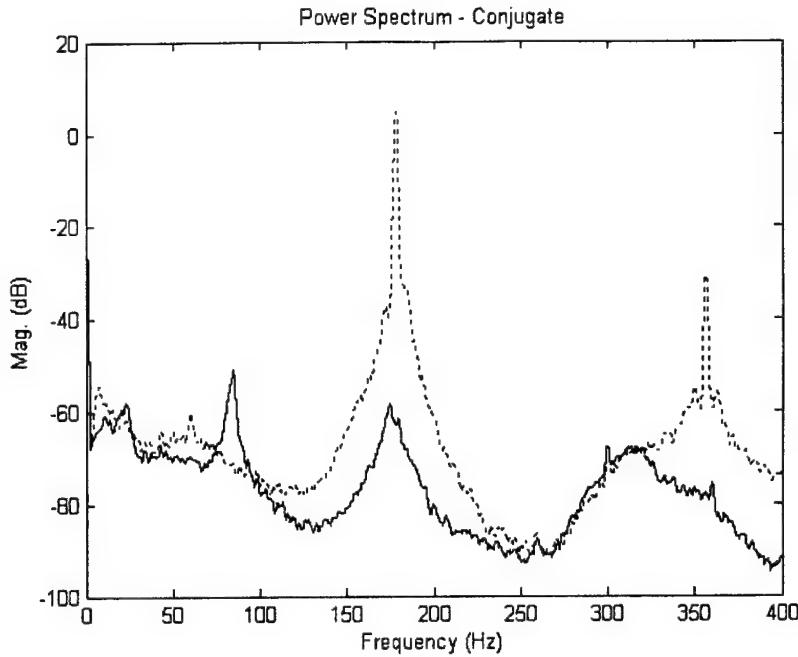


Figure 3.6 Power spectrum of Conjugate-controlled Rijke Tube, $\Phi = 0.544$

Results for the Conjugate algorithm are very similar to that of the Gradient algorithm, including the dominant energy component shifting to a sub-harmonic frequency.

The convergence times, ultimate attenuation, and the phase and magnitude of the controllers for each of the algorithms at this operating condition are tabulated in Table 3-1. For the low equivalence ratio case, each algorithm was successful in gaining control of the thermoacoustic instability. On average, there was about 62 dB of attenuation and a convergence time of approximately 1.4 seconds. Comparing the results from Table 3-1 with that of the pattern search algorithms of the last chapter demonstrates that the gradient descent algorithms outperformed the pattern search algorithms with respect to convergence time. The Rosenbrock and Hooke and Jeeves results showed that these algorithms might take steps in the non-optimal direction and thus produce temporary growth in the pressure oscillations during the convergence process, whereas the explicit gradient algorithms always move in the gradient direction towards the minimum.

Table 3-1: Convergence times and ultimate attenuation at limit cycle frequency for, $\Phi = 0.544$ Case

Algorithm	Convergence Time	Ultimate Attenuation	Filter Magnitude	Filter Phase
TAG	1.343 sec.	58.1 dB	0.84	112°
Gradient	1.363 sec.	66.2 dB	1.59	107°
Conjugate	1.478 sec.	63 dB	1.61	114°

3.4.1.2 Medium Equivalence Ratio Case

Raising the equivalence ratio to 0.582 did not negatively affect the performance of the algorithms. The results for this case, in fact, are very similar to that of the previous case. The results are shown in Figure 3.7 through Figure 2.20, encompassing each algorithm.

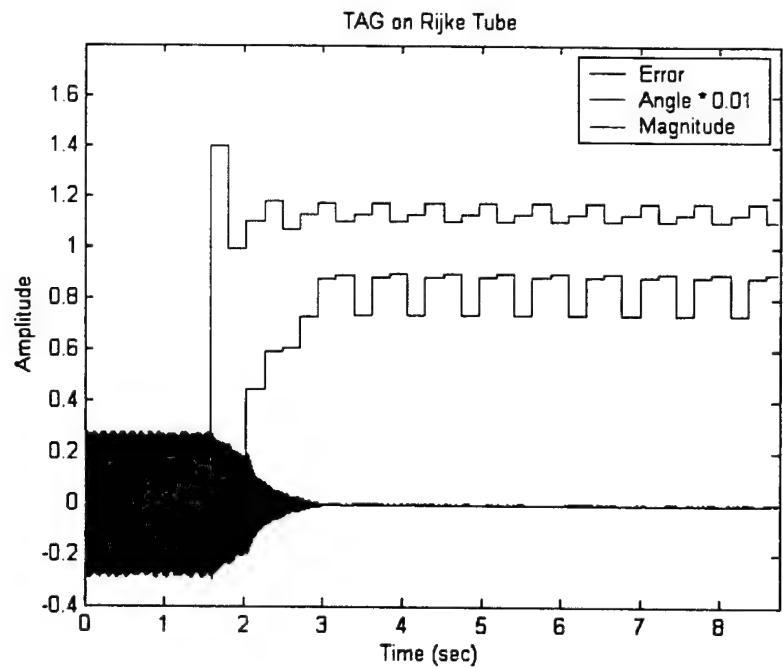


Figure 3.7 Phase and magnitude of filter TAG algorithm, $\Phi = 0.582$

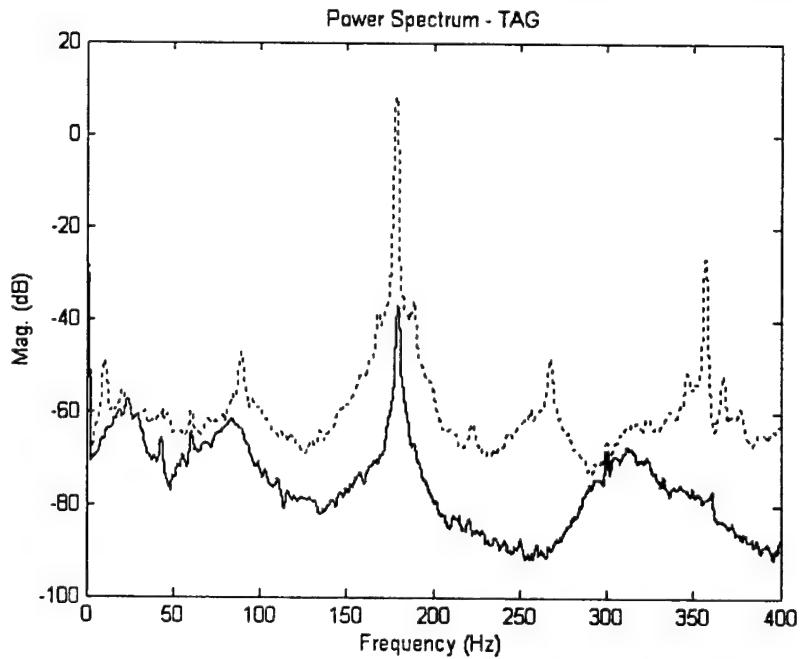


Figure 3.8 Power spectrum of TAG-controlled Rijke Tube, $\Phi = 0.582$

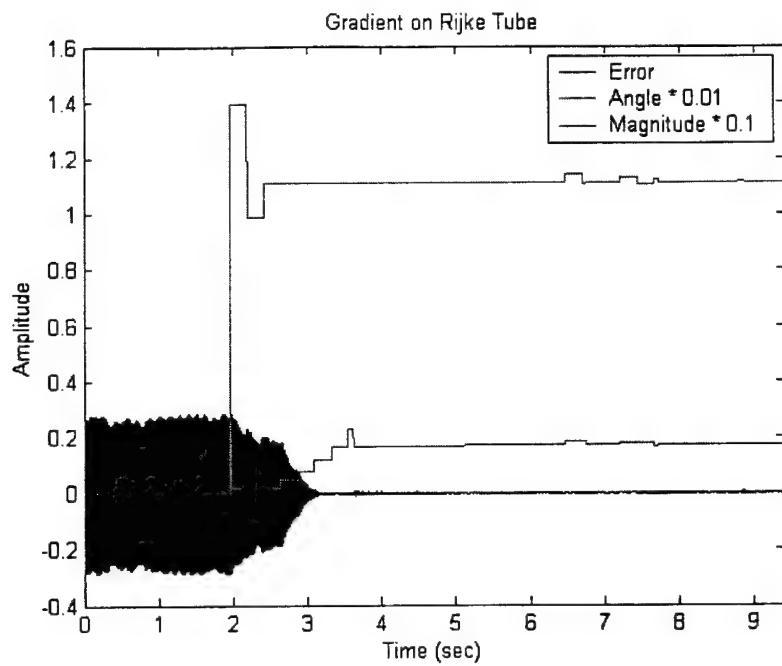


Figure 3.9 Phase and magnitude of filter Gradient algorithm, $\Phi = 0.582$

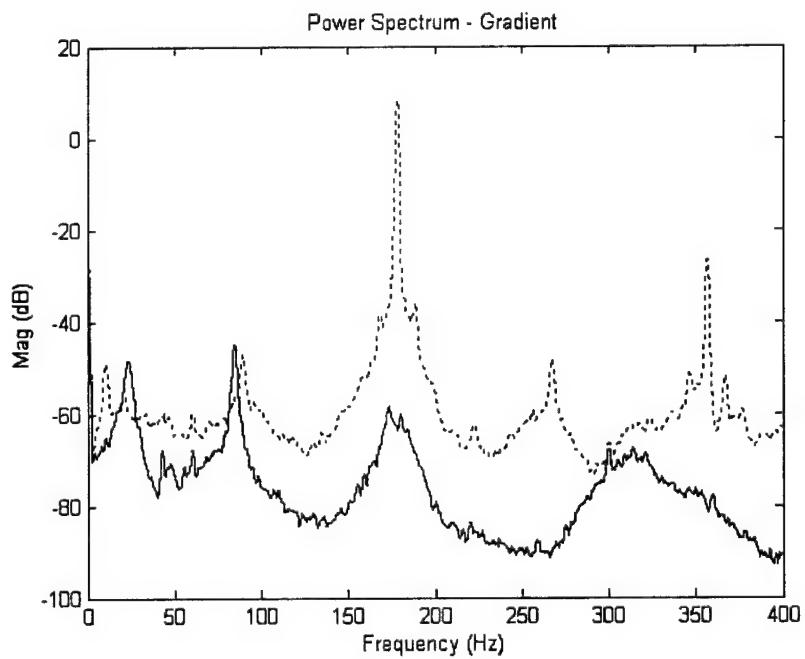


Figure 3.10 Power spectrum of Gradient-controlled Rijke Tube, $\Phi = 0.582$

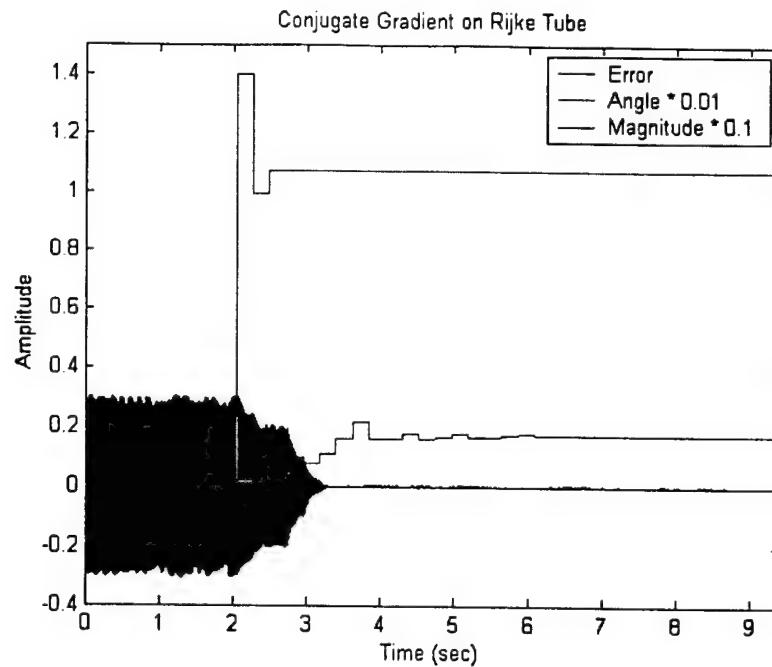


Figure 3.11 Phase and magnitude of filter Conjugate algorithm, $\Phi = 0.582$

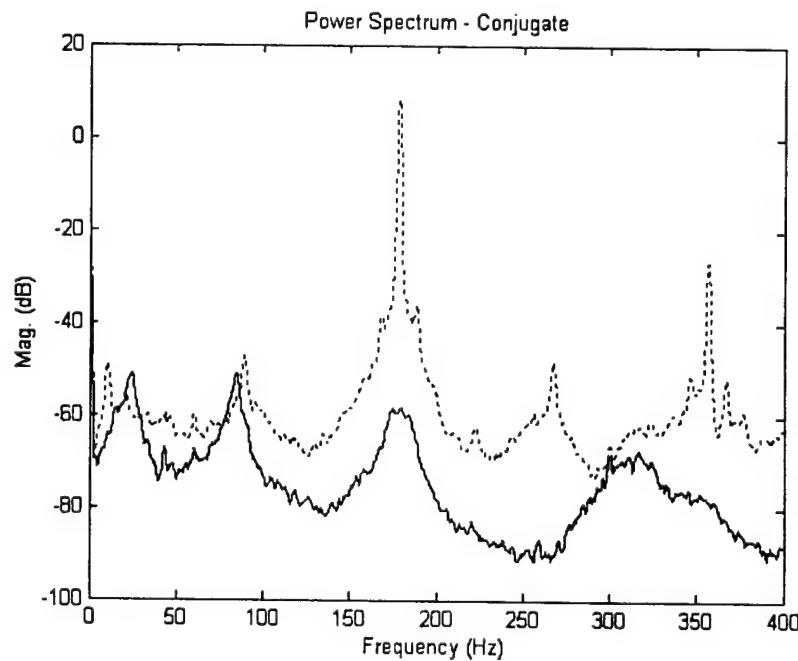


Figure 3.12 Power spectrum of Conjugate-controlled Rijke Tube, $\Phi = 0.582$

The results shown in the above figures are tabulated in Table 2-4, with the convergence time, ultimate attenuation, and magnitude and phase of the controller.

Table 2: Convergence times and ultimate attenuation at limit cycle frequency for the Medium Φ Case

Algorithm	Convergence Time	Ultimate Attenuation	Magnitude	Phase
TAG	1.636 sec.	44.8 dB	0.81	114°
Gradient	1.407 sec.	66.1 dB	1.68	111°
Conjugate	1.315 sec.	65.8 dB	1.69	107°

The magnitudes and phases that were developed by the algorithms for this operating condition are similar to those for the previous condition. The average convergence time for this case was about 1.45 seconds.

3.4.1.3 High Equivalence Ratio Case

Results for the TAG algorithm on the high equivalence ratio condition are presented in Figure 3.13 and Figure 3.14. Note that the explicit gradient algorithms, like the pattern search algorithms, are also affected by the lack of authority issue for this equivalence ratio.

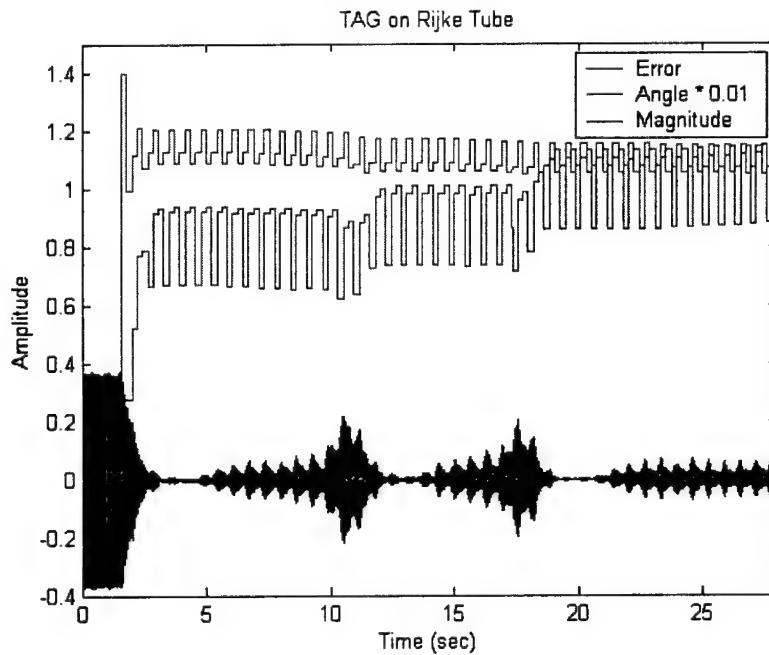


Figure 3.13 Phase and magnitude of filter TAG algorithm, $\Phi = 0.641$

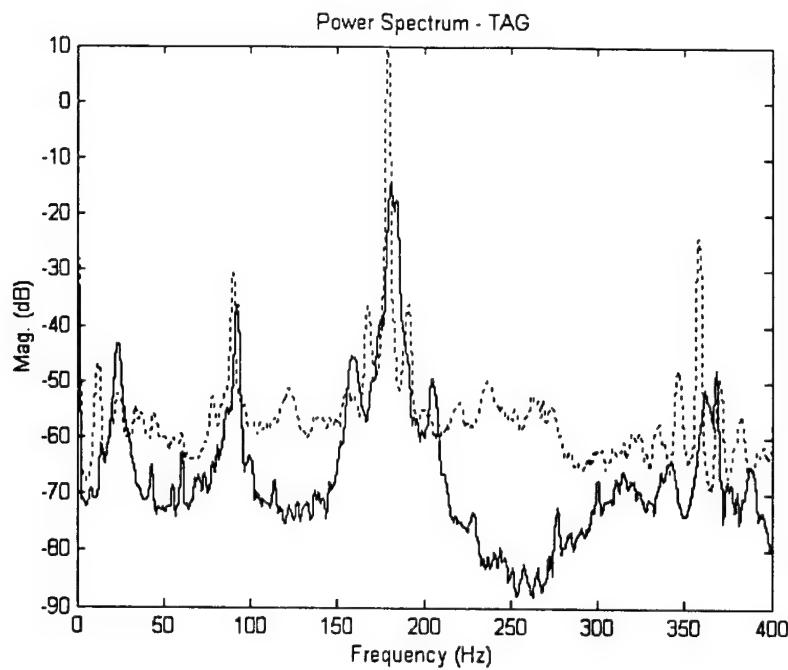


Figure 3.14 Power spectrum of TAG-controlled Rijke Tube, $\Phi = 0.641$

The results for the Gradient case are shown in Figure 3.15 and Figure 3.16, which also demonstrate a bursting behavior, though less severe than the TAG case.

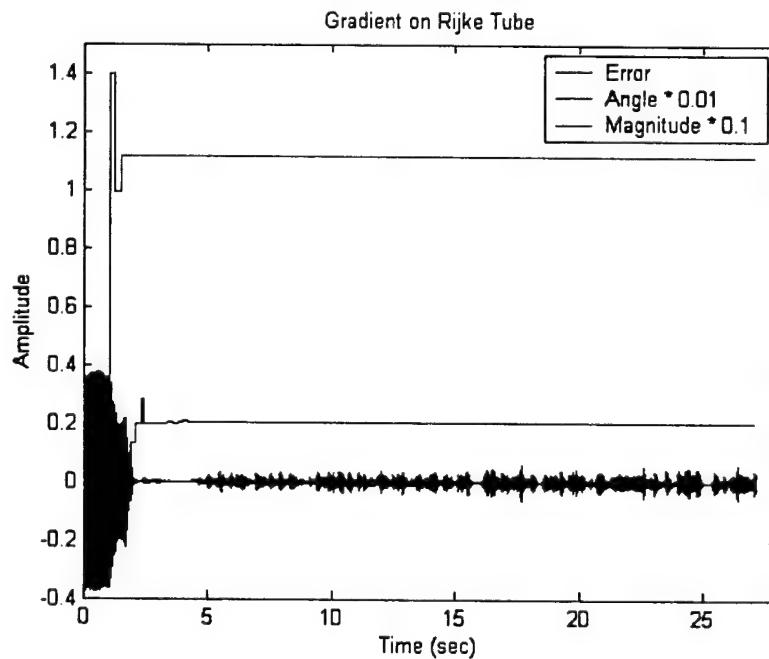


Figure 3.15 Phase and magnitude of filter Gradient algorithm, $\Phi = 0.641$

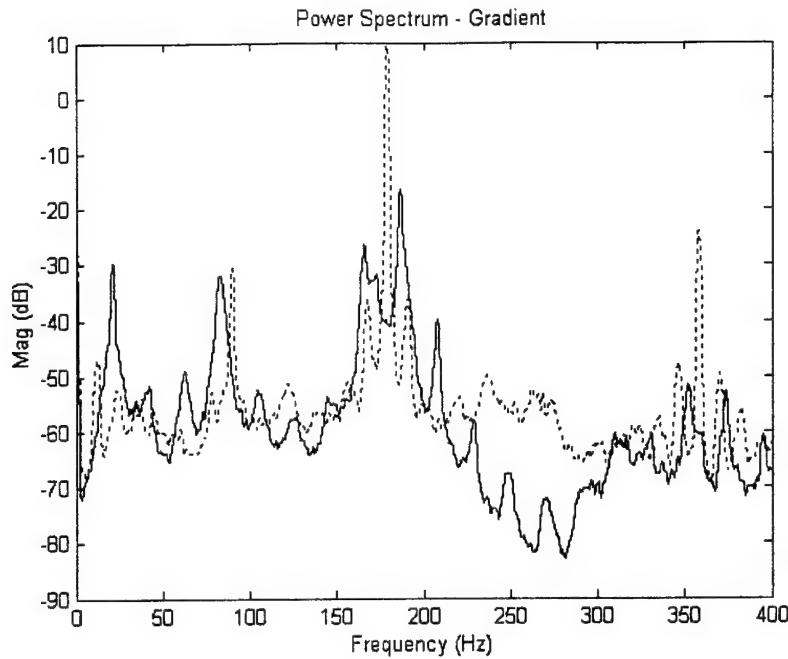


Figure 3.16 Power spectrum of Gradient-controlled Rijke Tube, $\Phi = 0.641$

Figure 3.17 and Figure 3.18 show the transient response and the power spectral density, respectively, for the Conjugate algorithm. The results of this test are very similar to that of the Gradient algorithm.

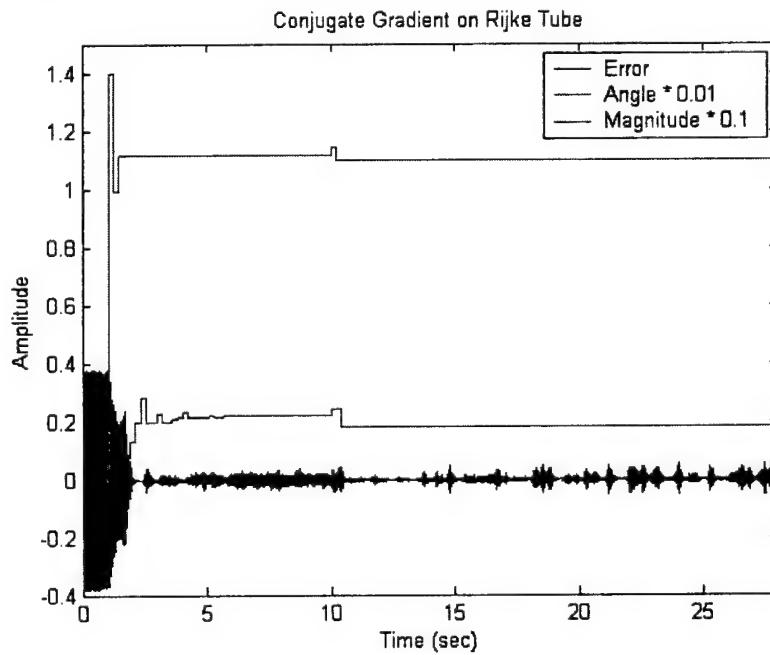


Figure 3.17 Phase and magnitude of filter Conjugate algorithm, $\Phi = 0.641$

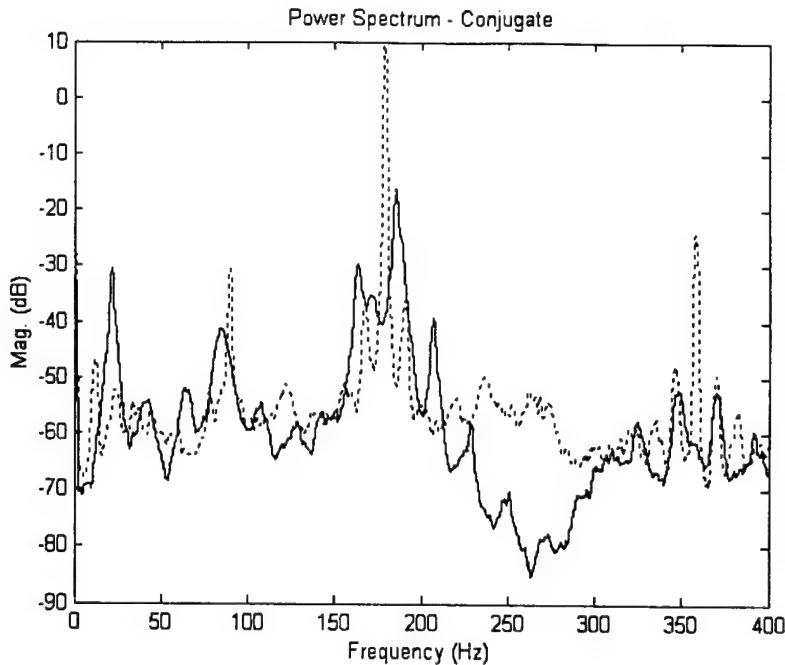


Figure 3.18 Power spectrum of Conjugate-controlled Rijke Tube, $\Phi = 0.641$

Much like pattern search algorithms at this high equivalence ratio, the explicit gradient algorithms could not maintain control over the Rijke tube instability. The performances of the explicit algorithms, however, were superior to that of the pattern search algorithms. The pattern search algorithms search the weight space and settles into a region that produces the lowest MSE of the pressure signal. The algorithms will begin to search again if a measurable increase in the MSE occurs. This is the reason that the bursts for these algorithms are so severe. Instead of continually computing the gradient like the explicit gradient algorithms, and thus gaining knowledge as the system changes, the pattern search algorithms wait until the bursts occur and then descend upon the optimal answer again.

3.4.2 Proportional Actuation Discussion

The three explicit gradient algorithms discussed here have been shown to achieve high levels of attenuation with good convergence rates when the system was within the authority range of the actuator. The Gradient algorithm and the Conjugate algorithm performed very similarly and, in general, better than the TAG algorithm. This is to be expected because of the incorporation of a line search and conjugate search directions in the Gradient and Conjugate algorithms, respectively.

When actuator authority was an issue, the algorithms converged but could not maintain an acceptable level of control. The explicit gradient algorithms did, however, outperform the pattern search algorithms discussed in the previous chapter.

3.5 References

- [1] M. A. Vaudrey. "Active Control Methods for Non-linear Self-excited Systems," PhD Dissertation, Virginia Tech, 2001.
- [2] M. S. Bazaraa, C. M. Shetty, *Nonlinear Programming: Theory and Algorithms*, John Wiley and Sons, Inc., New York, NY, 1979.
- [3] B. Widrow, S. D. Stearns, *Adaptive Signal Processing*, Prentice Hall, Inc., Englewood Cliffs, NJ, 1985.
- [4] D. M. Himmelblau, *Applied Nonlinear Programming*, McGraw-Hill, Inc., New York, NY, 1972.
- [5] K. McManus, T. Poinsot, S.M. Candel, "A Review of Active Control of Combustion Instabilities," *Prog. Energy Combust. Sci.* 19, 1993, pp. 1-29.
- [6] R.L. Raun, M.W. Beckstead, J.C. Finlinson, K.P. Brooks, "A Review of Rijke Tubes, Rijke Burners and Related Devices," *Prog. Energy Combust. Sci.* 19, 1993, pp. 313 – 364

4 Filtered-E Based Algorithms

4.1 Introduction

In this chapter, we explore the applicability of adaptive signal processing methods to the problem of active combustion control (ACC). These adaptive filtering methods, such as the LMS algorithm proposed by Widrow and Stearns (1985), have been highly successful in the active noise and vibration control community for narrowband disturbance rejection. The most relevant application of LMS control for noise or vibration control was the so-called Filtered-U algorithm [14] which required feedback compensation for the reference sensor. The ACC implementation is similar in its need for reference sensor compensation because the reference sensor is often identical to the error sensor. A number of combustion researchers have investigated the usefulness of these LMS algorithms for ACC ([1], [2], [11], [16], and [13]). However, most of the results from those experiments and numerical simulations indicated an uncertainty about the proper implementation with regard to employing a plant estimate. As a result, the literature shows that the LMS controlled systems often diverged, sometimes minutes after the combustor's pressure signature seemed to have been reduced to acceptable levels.

This chapter extends the existing research into the design of LMS adaptive filters for suppression of combustor thermoacoustic instabilities. Thermoacoustic instability control is fundamentally different from the noise and vibration control problems that have been widely discussed in the literature. The most notable differences are that there is no exogenous disturbance to be cancelled and the homogenous system to be controlled is both unstable and nonlinear. The specific effects that these differences have on the LMS adaptive control performance have been largely glossed over in earlier references. The primary objective of this chapter is to clarify the precise nature of the LMS control performance for thermoacoustic instability control applications. In particular, we propose a completely different alternative for the system model used in the filtered-X adaptive feedback structure. A system identification *method* to estimate the required model is also proposed and it is shown that this method leads to a unique linear approximation that can be used for stabilizing control. Extensive experiments are also conducted on a Rijke tube combustor, showing indefinite stabilizing control of the limit cycle oscillation. Nonlinear simulations of relevant ACC LMS control experiments are used to explain the detailed behavior of the plant for those cases where control is achieved. Then we examine in detail an operating condition for which the LMS controller is not effective and the system exhibits bursting behavior. It is shown that this behavior is unrelated to the adaptive controller and is due, ultimately, to a lack of control authority. (We note that bursting is a characteristic of certain adaptive controllers and that bursting has been noticed for both fixed-gain and adaptive ACC control by other researchers.)

Although our experiments with LMS-based feedback control have been successful, it is important to consider potential problems with the algorithm. LMS-based feedforward algorithms have substantial stability guarantees, but these do not carry over to feedback implementations. Thus, we examine the potential for feedback loop instabilities and algorithm divergence in terms of the plant estimate for LMS-based feedback control of stable systems subject to exogenous disturbances. Next, the results are extended to filtered-U control structures as well as the control of unstable, self-excited plants. In addition, the

dependence of the instabilities on each other, and the use of online system identification techniques are examined for both adaptive feedback and filtered-U control. Finally, two simulations are presented. The first illustrates a feedback loop instability independent of algorithm divergence. The second simulation shows how a conventionally accurate plant estimate can yield an algorithm divergence independent of a loop instability. This is shown to agree with the analytical results.

4.2 Adaptive Feedback Control

4.2.1 Classical Disturbance Suppression

The adaptive feedback controller considered in this chapter can be viewed as a special case of Filtered-U control where the control-to-reference path and the control-to-error path are identical because the error sensor also serves as the reference sensor [14]. Figure 4.1 shows the adaptive feedback control block diagram for the case of an external disturbance and stable plant. It is instructive to briefly examine this system as a precursor to analyzing LMS control of a self-excited thermoacoustic system.

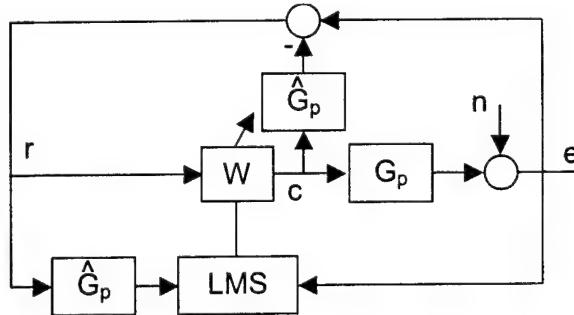


Figure 4.1 Adaptive Feedback External Disturbance

In a purely feedforward control system, a separate reference signal that is correlated with the disturbance is used as the input to the adaptive filter. It is well-known that correct estimation of the dynamic phase, within ninety degrees, is sufficient to prevent divergence of the LMS gradient search [7]. In adaptive feedback control, the reference signal is derived directly from the error sensor by removing the component of the error signal that is due to the control signal, leaving only a measure of the disturbance. If we let the control signal be the output of the FIR adaptive filter (W), we note that it must go through the plant dynamics (G_p) before acting on the external disturbance (n). The dynamics represented by the plant include everything present in the control signal to error signal path, including the A/D and D/A. Therefore, removing the control signal component from the error requires subtracting the output of the adaptive filter, filtered by the plant estimate, from the measured error signal. Hence, if the plant estimate is not perfect, a non-zero feedback path is introduced. The closed loop transfer function between the error and the disturbance can then be written as:

$$\frac{e}{n} = \frac{1 + \hat{G}_p W}{1 + \hat{G}_p W - G_p W} \quad (4.1)$$

Unlike the self-excited system, if the plant estimate is exact ($\hat{G}_p = G_p$), there is no feedback loop as the denominator of the above transfer function vanishes and the system behaves as a strictly feedforward system where \hat{G}_p and W are both stable systems. For the case where $\hat{G}_p \neq G_p$, it is clear that the poles of (4.1) will change with the adaptation of the filter W and represent a potential source of controller instability.

4.2.2 Self-Excited Systems

Figure 4.2 illustrates a simplified block diagram of an adaptive feedback controller applied to a self-excited system, such as the VPI Rijke tube combustor. Although the physical dynamics of the self-excited system are complex and the subject of ongoing research, a simplified model consisting of a linear acoustics block, G_A , a linear flame dynamics block, G_F , and a nonlinear coupling block is sufficient for the control system analysis discussed here. In the analysis and simulation, the nonlinear block is treated as a static, saturating-type nonlinearity. Though crude from a broadband modeling point of view, the extremely tonal nature of the system makes this a useful approximation.

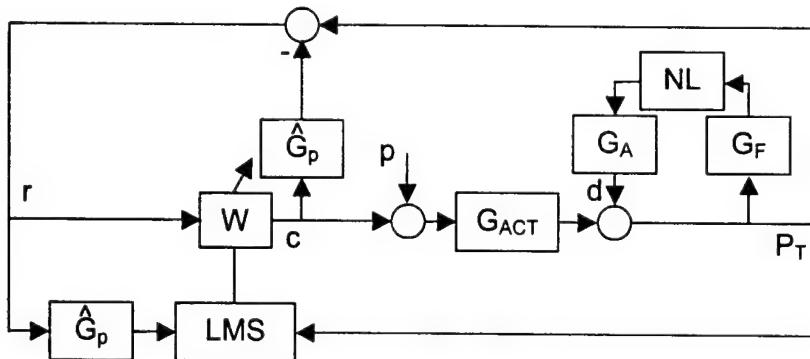


Figure 4.2 Adaptive Feedback Self-Excited

The actuation signal passes through a linear block, G_{ACT} , that incorporates both the actuator dynamics and the acoustic dynamics. Since the outputs of G_A and G_{ACT} may both be considered as acoustic pressures, they can be superposed to produce the total combustor pressure variation, P_T . In the controller, the output of the adaptive filter is filtered by the plant estimate and used to generate the reference signal from the error signal as before, although the appropriate plant estimate is less obvious now as a result of the inclusion of the self-excited system.

The controlled system consists of two main loops, the physical feedback loop and the control loop, as shown in Figure 4.3. We define the optimal controller as the controller that completely nullifies the physical feedback loop. Equating the top loop to the bottom loop (with a minus sign) and solving for the optimal adaptive filter (W_{OPT}), results in

$$W_{opt} = \frac{-G_A G_F}{G_{ACT} + G_A G_F \hat{G}_p} \quad (4.2)$$

By substituting (4.2) into the block diagram of Figure 4.3 as the adaptive filter W , it is obvious that the lower loop (feedback controller) exactly cancels the upper loop $G_A G_F$.

resulting in a completely stabilized closed loop system. Equation 4.3 illustrates this result for the probe input shown in Figure 4.3.

$$\frac{P_T}{p} = \frac{G_{ACT}(1 + \hat{G}_p W)}{(1 - G_A G_F)(1 + \hat{G}_p W) - G_{ACT} W} \Big|_{W=W_{opt}} = G_{ACT} \quad (4.3)$$

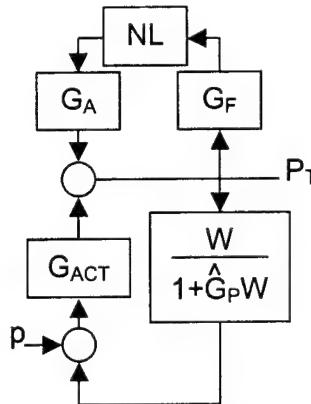


Figure 4.3 Redrawn Controlled System

Discussion of the optimal adaptive filter weights is more interesting for this self-excited plant application, versus the ANC problem, because of the quiescent state that the combustor will reach as soon as the poles have crossed back into the left-*half* Laplace plane under the action of the controller. Because the open-loop self-excited system of (4.3) does not contain dynamics that are on the imaginary axis (i.e. marginally stable), there is a set of gains (between the imaginary axis and the optimal solution) that will stabilize the system without requiring the adaptive filter to reach its optimal solution. This system will be more lightly damped than the open-loop system but still stabilized. If the adaptation causes the system to stabilize, the error signal will go to zero and the adaptation will stop, never reaching the optimal gain.

4.3 Analysis of Adaptive Feedback Applied to Self-Excited Systems

4.3.1 The Correct Plant Estimate

Initially, the system will be examined in the linear range so that the nonlinearity reduces to a simple gain that can be lumped with the flame dynamics, G_F . The total pressure (P_T) serves as the error signal to be reduced and is used by the adaptive feedback structure to update the weights and create the reference signal. The expression for the transfer function between the probe input of Figure 4.3 and P_T results in:

$$\frac{P_T}{p} = \frac{G_{ACT}(1 + \hat{G}_p W)}{1 - G_A G_F - G_A G_F \hat{G}_p W + \hat{G}_p W - G_{ACT} W} \quad (4.4)$$

In a purely feedforward situation, the dynamics from the control-to-error path are clear, and can be identified directly, often in the absence of the disturbance, and a model can be generated. The appropriate choice for the control-to-error path estimate is not as clear when the self-excited system is addressed. The self-excited system of Figure 4.3 offers two

logical choices for the plant estimate. Following the standard procedure for feedforward problems, one choice for the plant estimate is the open loop dynamics of the control-to-error path, which yields

$$\frac{P_T}{p} = \hat{G}_p = \frac{G_{ACT}}{1 - G_A G_F} \quad (4.5)$$

The acoustic dynamics have been moved from the *forward* path of the self excited loop and appear separately as part of the actuator dynamics and the self-excited *feedback* loop. (Stated otherwise, the same acoustics influence the self-excited loop and the actuator path).

Substituting this choice for the plant estimate, (4.5), into the denominator of (4.4) yields

$$1 - G_A G_F - G_A G_F \hat{G}_p W + \hat{G}_p W - G_{ACT} W \Big|_{\hat{G}_p = \frac{G_{ACT}}{1 - G_A G_F}} = 1 - G_A G_F \quad (4.6)$$

This result guarantees that the adaptive filter cannot influence the poles of the closed loop system, and they will remain in the unstable right half plane! Therefore, this choice of the plant estimate is not considered to be a valid solution and can never robustly stabilize the self-excited system.

An alternative proposal is to use the actuator path alone as the plant estimate such that $\hat{G}_p = G_{ACT}$. Assume the output of the acoustic portion of the self-excited loop is the external disturbance to be canceled at the error sensor. In this case, the dynamics between the control output and the error sensor represent the actual control-to-error path. In addition, the artificial reference signal now becomes an estimate of the exact disturbance signal at the error sensor. This can be seen by recognizing that:

$$\begin{aligned} P_T &= d + G_{act} c \\ r &= P_T - \hat{G}_p c = d + G_{act} c - \hat{G}_p c \end{aligned} \quad (4.7)$$

where d is the output of the acoustic plant, c is the output of the controller, P_T is the measured error signal and r is the derived reference signal. Not only is the reference signal accurately representing the disturbance to be cancelled (when the plant estimate is G_{ACT}), but the required plant estimate is of a strictly stable system. Using G_{ACT} as the control-to-error path estimate, we can again examine the closed loop system of (4.4).

$$\frac{P_T}{p} = \frac{G_{ACT} (1 + \hat{G}_p W)}{1 - G_A G_F - G_A G_F \hat{G}_p W + \hat{G}_p W - G_{ACT} W} \Big|_{\hat{G}_p = G_{act}} = \frac{G_{ACT} (1 + G_{ACT} W)}{1 - G_A G_F - G_A G_F G_{ACT} W} \quad (4.8)$$

It is clear from (4.8) that the adaptive filter can now influence both the zeros and the poles of the closed loop control system, allowing for the possibility of stabilizing the self-excited system.

If we assume that the plant estimation error is zero and \hat{G}_p is exactly equal to G_{ACT} , the block diagram of Figure 4.2 can be redrawn as shown in Figure 4.4. It is clear that when using the appropriate plant estimate (assuming perfect identification), the system is identical to the feedforward filtered-X structure and can theoretically behave as a feedforward system. In view of Figure 4.3, it is easy to see how the adaptive filter imparts the needed gain and phase to control the limit cycling system through the feedback loop. The significance of the control-to-error path estimate (\hat{G}_p) in the controller transfer function is also evident from Figure 4.3.

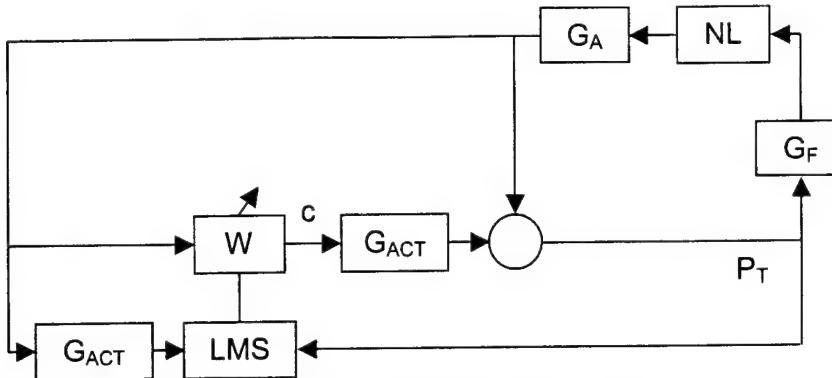


Figure 4.4 Perfect Identification Diagram Redrawn

4.3.2 System Identification

Because of the uncertainty in previous work concerning the plant to be estimated, the topic of system identification has not been addressed adequately. Billoud et.al. [2] discuss using the filtered-X algorithm where the plant estimate is either generated from a single tone identification near the limit cycle frequency or by using an assigned gain and phase delay. Kemal and Bowman [11] discuss obtaining an open loop frequency response but do not provide details regarding the exact nature of the transfer function obtained.

Ultimately, they note that using a gain and delay appears sufficient for tonal control.

Koshigoe et. al. [13] among others, have investigated using an online system identification procedure. Because of the feedback nature of the self-excited system, the online identification will necessarily include the controller and feedback dynamics. If G_{ACT} is the correct estimate, online identification procedures cannot produce the correct result after the loop has been closed.

The identification problem, then, is to find a stable representation of G_{ACT} when we have access to the input to G_{ACT} but the only measurable output is P_T . Extinguishing the flame so that the physical feedback loop disappears is not an option since the hot acoustics, which are an integral part of G_{ACT} , are very different from the cold acoustics and are directly influenced by the flame temperature and temperature gradient. Thus, the identification must take place in the presence of the thermoacoustic limit cycle.

A new approach to identifying the open-loop plant relies on a probe signal consisting of low-amplitude sinusoids at frequencies near the limit-cycle frequency and within the passband of any bandpass filters used to filter the pressure signal before control. Using a Fourier transform of the output signal, the frequency response of the plant at a small number of frequencies can be determined. Using a least-squares approach, a low-order, discrete-time model can be fit to this data and used as the model of the stable plant.

Since the open-loop plant is in a steady limit cycle, it is not immediately obvious whether such an identification approach will produce a reasonable linear model or whether such a model will be stable or unstable. Our experimental results have shown that very low-order linear models can account accurately for the frequency response data and that these models are always stable.

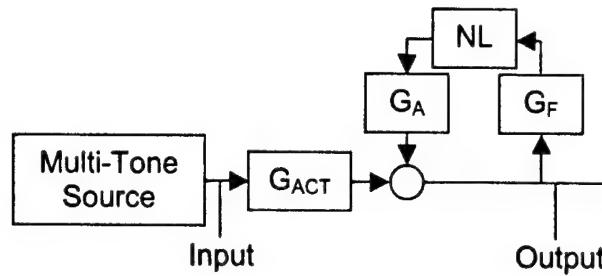


Figure 4.5 System Identification – Open Loop

To understand this, consider the block diagram in Figure 4.5. From a describing function analysis, the gain of the limit cycle through the static nonlinearity is such that the total gain around the loop is unity. When a probe signal is injected into the system, we expect that the frequency response at the probe frequency will be that of the linear system made up of the linear blocks in the diagram and with the nonlinear block replaced by a suitable linear gain.

To determine the value of this gain, consider an input to the nonlinearity of the form

$$x = A_1 \sin \omega_1 t + A_2 \sin(\omega_2 t + \theta) \quad (4.9)$$

where ω_1 is the frequency of the limit cycle, ω_2 is the frequency of the probe and $A_1 \gg A_2$. A first order approximation of the output of the nonlinearity, $f(x)$, is given by

$$f(A_1 \sin \omega_1 t + A_2 \sin(\omega_2 t + \theta)) = f(A_1 \sin \omega_1 t) + f'(A_1 \sin \omega_1 t) A_2 \sin(\omega_2 t + \theta) \quad (4.10)$$

and will be valid for A_2 sufficiently small. The gain of the nonlinearity at the limit cycle frequency is given by

$$\frac{2}{A_1 T} \int_0^T [f(A_1 \sin \omega_1 t) + f'(A_1 \sin \omega_1 t) A_2 \sin(\omega_2 t + \theta)] \sin \omega_1 t dt \quad (4.11)$$

where T is the length of a period of the overall waveform. If no period exists, then the limit as $T \rightarrow \infty$ can be taken. By arguing that the integral of incommensurate frequencies will vanish, the second term in the integral disappears and the linear gain associated with the limit cycle frequency is equal to the describing function of the nonlinearity. That gain is written as:

$$g_{LC} = \frac{2}{A_1 T} \int_0^T f(A_1 \sin \omega_1 t) \sin \omega_1 t dt \quad (4.12)$$

The gain of the nonlinearity at the probe frequency is given by

$$\frac{2}{A_2 T} \int_0^T [f(A_1 \sin \omega_1 t) + f'(A_1 \sin \omega_1 t) A_2 \sin(\omega_2 t + \theta)] \sin(\omega_2 t + \theta) dt \quad (4.13)$$

Arguing as before, the first term in the integral will go to zero. In addition, the only part of the second term that will contribute to the integral is the constant component of $f'(A_1 \sin \omega_1 t)$ times the constant component of $A_2 \sin^2(\omega_2 t + \theta)$. Thus, the gain at the probe frequency is given by

$$g_{Pr} = \frac{1}{T} \int_0^T f'(A_l \sin \omega_l t) dt \quad (4.14)$$

Note that the gain of the probe signal is independent of the amplitude and frequency of the probe signal, subject to the restriction that the amplitude of the probe is small. For the tanh nonlinearity considered in this section, the limit-cycle and probe gains can be computed numerically and are shown in Figure 4.6. The gain of the probe signal is less than the gain of the limit cycle signal through the nonlinearity. Since the gain of the limit cycle frequency is just that value needed to make the closed-loop system marginally stable, the lower probe gain will cause the closed-loop system to appear stable. Since the probe gain is not a function of the probe frequency, the system identified by considering the frequency response at the probe frequency will appear linear and stable.

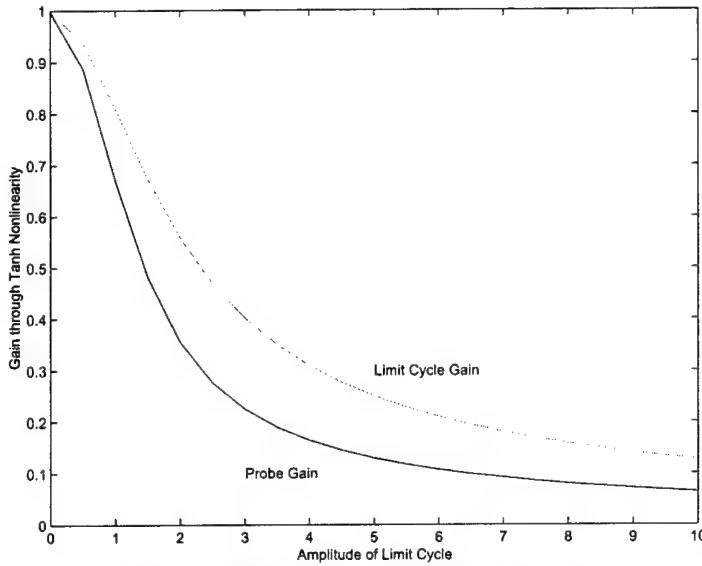


Figure 4.6 Probe Gain and Limit Cycle Gain

Rewriting 4.5 shows that the identified system will have the transfer function

$$\frac{P_T}{p} = \hat{G}_p = \frac{G_{ACT}}{1 - G_A G_F g_{Pr}} \approx G_{ACT} \quad (4.15)$$

The reason that this transfer function is approximately equal to G_{ACT} over the bandwidth of interest is because the denominator is very close to one at the probe frequencies. There are two reasons for this. First, the factor $G_A G_F g_{Pr}$ is always less than unity for all frequencies in the bandwidth. Secondly, G_A contains lightly-damped acoustic poles that will have a high gain very near the instability frequency. The gain will be significantly lower a small distance from this frequency (see Figure 4.10), where the probing is actually performed. Thus, at the probe frequencies, the factor $G_A G_F g_{Pr}$ will be significantly less than unity, resulting in accurate measurements of G_{ACT} .

This plant estimate is stable at all probe frequencies, thereby avoiding the issue of attempting to use an unstable system identification, or a time-delay plant, as other investigations have discussed. In addition, it approximates G_{ACT} , previously shown to be the desired transfer function of the plant estimate. On a related note, we point out that this

system ID method also makes it possible to design a feedback controller a priori that will apply the correct phase to the system while avoiding secondary peaks induced from the controller feedback loop (Vaudrey et al., 2000).

4.4 Simulation, Experimental Results, and Actuator Authority

4.4.1 Simulation Description

A simulation has been designed to permit rapid and easy investigation of the performance of the adaptive feedback controller in conjunction with the self-excited, limit cycling system. As will be seen, results from the simulation agree closely with the behavior of the actual experiment. This is undoubtedly due to the fact that the dynamics that dominate the system are a single pair of lightly-damped acoustic poles.

Figure 4.2 represents the general form of the adaptively controlled self-excited system that was simulated. The self-excited loop consists of a low pass filter to represent a model of the flame dynamics (G_F), a static nonlinearity represented by a hyperbolic tangent function, and a single-mode lightly-damped (2% viscous damping) acoustical model (G_A) at approximately 175 Hz. The loop gain and nonlinearity gain were adjusted to yield a steady limit cycle after approximately 2 seconds at a sampling rate of 1600 Hz. The actuator path consists of the same acoustical model plus some amount of phase delay. For the experiment, this delay represents all components in the control-to-error path.

The simulation runs in two separate modes, as does the experimental setup. After the limit cycle is established, a multi-tone probe signal is applied to the open loop plant as shown in Figure 4.5. The input to output relationship at the probe frequencies establishes the magnitude and phase of the linear plant as described in the system identification section above and in (4.15). This discrete frequency response data representing G_{ACT} is then used to generate a least squares infinite impulse response (IIR) transfer function fit in the z-domain, typically of order less than 6, with a pole very near the unit circle representing the "hot" acoustic mode. This fit is then used as the plant estimate, G_{ACT} , during the second mode of the simulation. After the limit cycle has reached a steady state, and the probe frequencies have been turned off, the plant model mentioned above is used in the adaptive feedback control loop shown in Figure 4.2, both as the filtered-X part of the LMS algorithm and as the plant estimate used to derive the reference signal.

The simulation illustrated here is a case having a relatively low heat release, corresponding to a low gain in the self-excited feedback loop. Since broadband control is not the goal, only two adaptive filter weights were used to control the single tone instability. In the experimental setup, a steep (8-zero, 16-pole) bandpass filter was used to eliminate frequency content other than the limit cycle tonal. In the simulation, the only significant content in the error signal is the limit cycle sinusoid so a bandpass filter was not necessary. Using the stable plant model of G_{ACT} generated from the fit of the FRF obtained from the process shown in Figure 4.5, and a relatively fast convergence parameter for the LMS weight update equation, it is seen that the two adaptive weights shown in Figure 4.7 quickly reach a steady state condition that completely stabilizes the limit cycle as shown by the controlled and uncontrolled (dotted) power spectra in Figure 4.8.

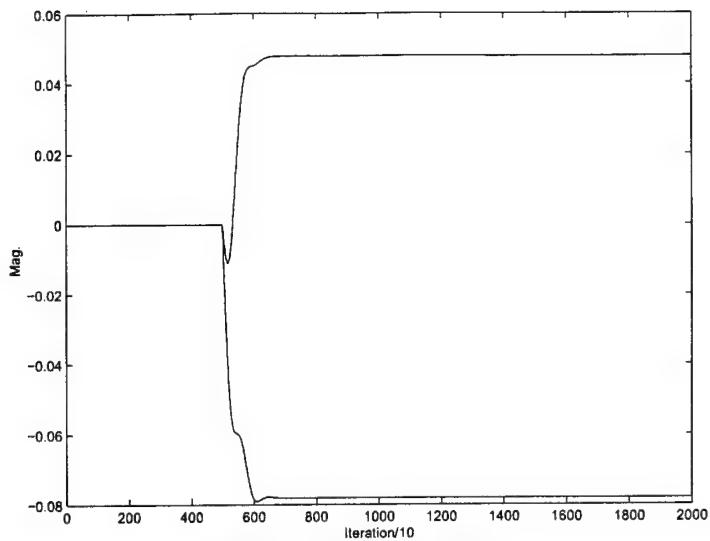


Figure 4.7 Weights in Time

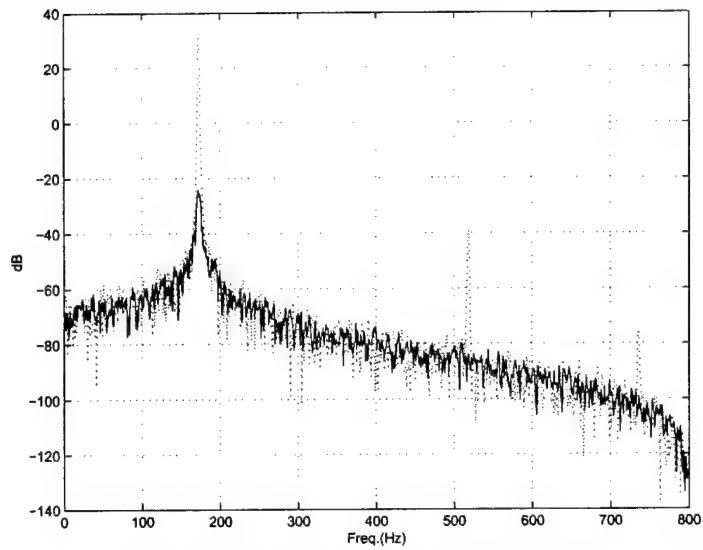


Figure 4.8 Low Heat Release LMS Simulation

As discussed in the analysis section above, it is not required that the adaptive filter achieve the optimal gain in order to drive the error signal to zero. Figure 4.9 clearly illustrates this by showing the path of the magnitude and phase of the actual adaptive filter during adaptation, as compared to the optimal magnitude and phase as computed from (4.2).

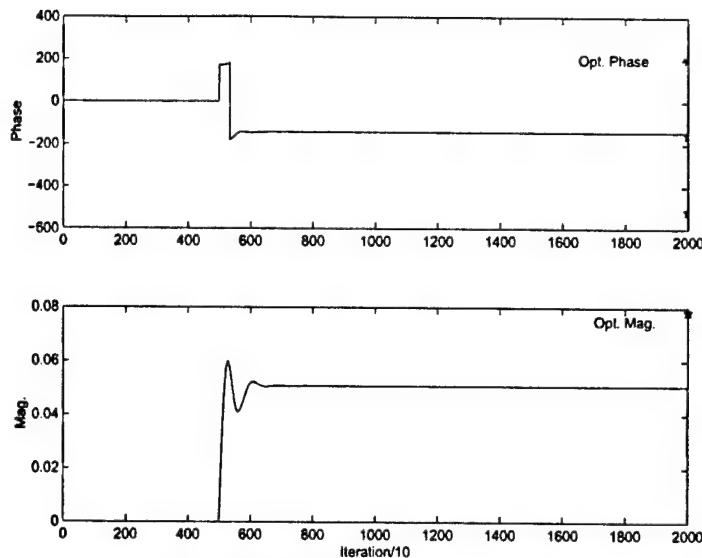


Figure 4.9 Low Heat Release LMS Path to Optimal

4.4.2 Experimental Results

The experimental process was identical to that described in the simulation section above. A multi-tone FRF was performed on the entire plant and a least squares fit was applied to the data to generate an IIR filter model that represented the control-to-error path. This model was then used in the structure shown in Figure 4.2.

With the heat release at a relatively low gain (controlled by adjusting the premixed methane equivalence ratio to a value of 0.51 and a total flow rate of 120 cc/sec), the Rijke tube combustor instability at 175 Hz was stabilized indefinitely with a two weight adaptive filter and the stable plant model obtained prior to control. Figure 4.10 shows (in asterisks) the actual magnitude and phase data collected from the tube along with the 6th order model frequency response shown as the solid line.

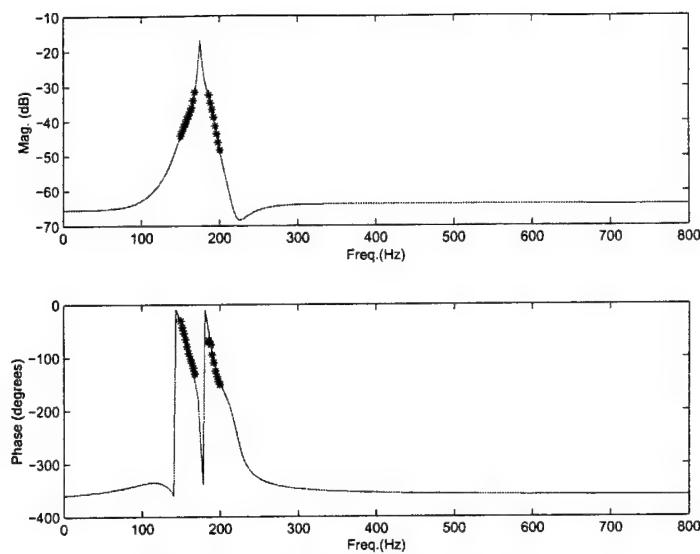


Figure 4.10 Experimental Model Data and Fit

Figure 4.11 shows the uncontrolled (dotted) and controlled power spectra of the total pressure in the tube (converted from voltage units). The second harmonic at 350 Hz disappears under control, revealing the shape of the third acoustic mode of the tube. The natural damping of the second acoustic mode is greater than that which is shown in Figure 4.11, but as discussed earlier, the system can be stabilized by having a pole in the left half plane that is more lightly damped than the natural acoustic mode.

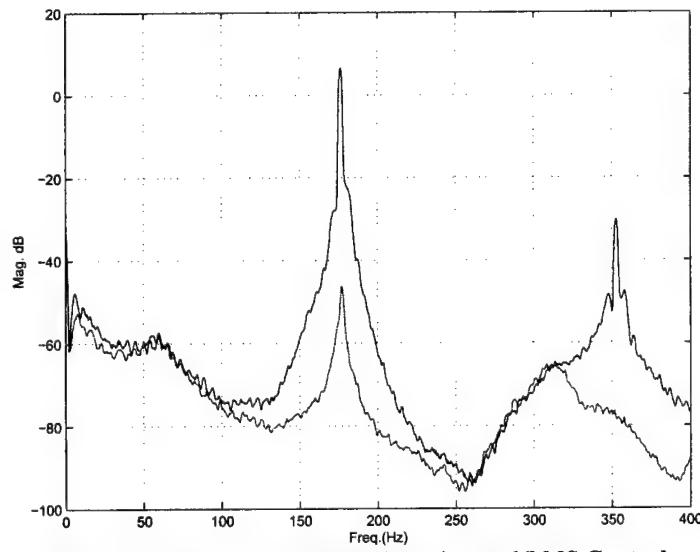


Figure 4.11 Low Heat Release Experimental LMS Control

A manually adjustable gain and phase shift controller was also applied to the same limit cycling system. Referring to Figure 4.3, it is apparent that the fixed gain controller replaces the transfer function:

$$\frac{W}{1 + \hat{G}_p W} \quad (4.16)$$

Examining the magnitude and phase of the converged Equation 16 at 175 Hz as compared to the magnitude and phase of the fixed feedback controller, it is seen in Figure 4.12 that the phase of both controllers is nearly the same at 175 Hz. The magnitude, however, is significantly lower for the adaptive system and will not increase with time because the error signal has been driven below the 1-bit noise floor of the A/D. It is known that excessive gain can produce controller-induced instabilities (Vaudrey et. al. 2000, Saunders et. al. 1999-2). However, the adaptive controller can (and does) prevent controller-induced instabilities by changing its shape and adjusting its magnitude to minimize the mean squared error.

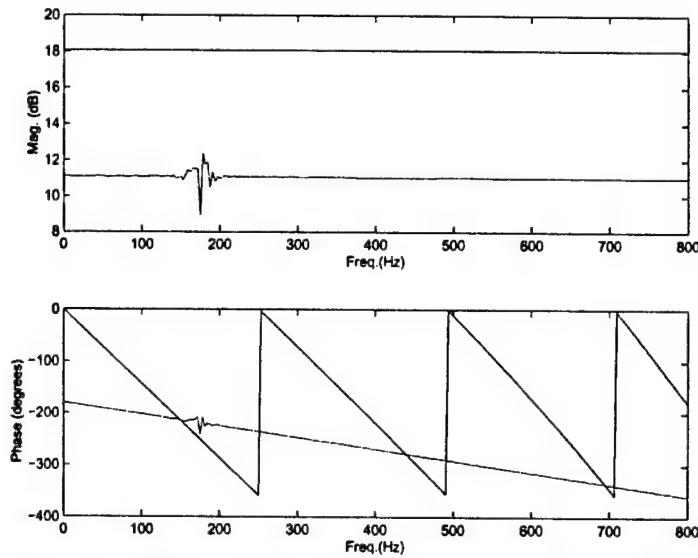


Figure 4.12 Fixed Gain and Adaptive Controller Comparison

The results discussed above are indicative of the results obtained at a number of operating conditions at low equivalence ratios. As the equivalence ratio, and hence heat release, was increased, a point was reached where intermittent behavior was observed. Figure 4.13 compares the convergence behavior in the time domain of the adaptive feedback LMS algorithm for low and high heat release conditions. The upper time trace represents the convergence corresponding to the control performance in Figure 4.11 whereas the lower trace exhibits the searching behavior present at the higher heat release operating conditions.

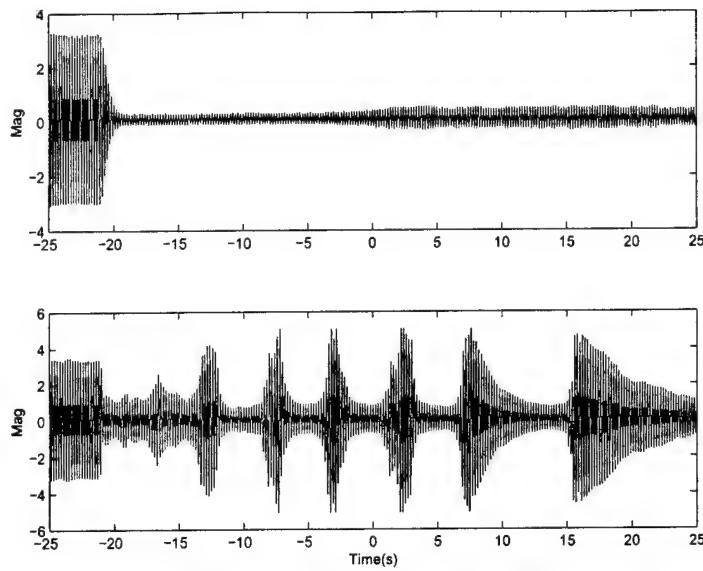


Figure 4.13 Low and High Heat Release LMS Control

To better understand the intermittency phenomenon, a fixed gain feedback controller was applied to the combustor at the same operating conditions. This removed the variability inherent in the time varying (adaptive) controller. As shown in Figure 4.14, the fixed gain, controlled system exhibited the same cyclic intermittency. Since the fixed gain feedback controller has no time varying components, it was determined that the heat release dynamics must be changing in a periodic manner to cause the intermittent behavior.

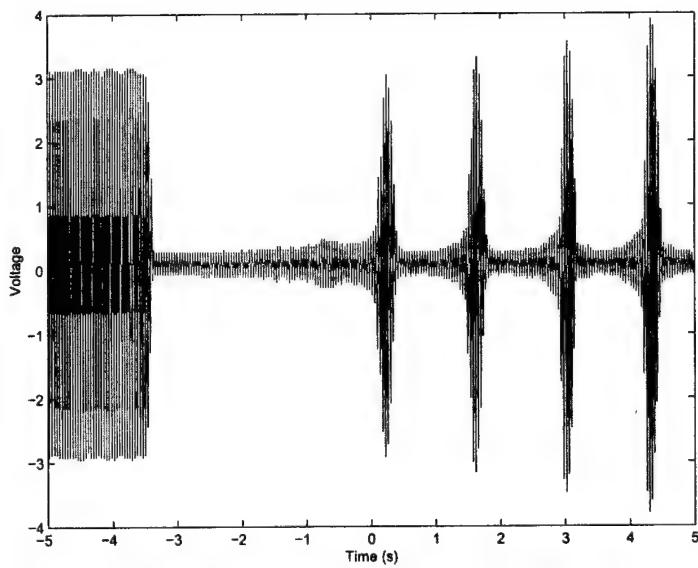


Figure 4.14 Fixed Gain Feedback Control – High Heat Release

4.4.3 Actuator Authority

Researchers have frequently witnessed intermittent behavior of self-excited systems when controlled by adaptive approaches. Quite often this behavior is attributed to algorithmic anomalies or controller induced instabilities. While each of these problems can contribute to searching or marginally stable behavior (Vaudrey, et al., 2001), a third reason for intermittent behavior is presented here which is not related to time-varying control. It is important to distinguish this form of intermittency from other instability mechanisms because the controller design itself is not at fault.

The self-excited combustion system is unique in that the plant dynamics change under the action of the controller. For a given set of limit cycling *operating conditions* described by a fixed equivalence ratio and flow rate, the combustor *operating point* can be described by the dynamic gain of the heat release that is a function of temperature, heat transfer rates to the combustor walls, and other unmodeled, low frequency dynamics. The pressure fluctuations from the limit cycle influence these low frequency dynamics. Consequently, when the pressure fluctuations are reduced by the feedback control, the operating point changes. This new, controlled, operating point can have enough increased heat release gain so that the control effectiveness is limited by actuator power. This effect can be more detrimental at higher heat release conditions where actuator authority limitations become relevant and required controller gain is high.

Based on the above observation, a hypothesis was developed and proven experimentally to define the searching phenomenon seen for fixed gain feedback control. The hypothesis can be explained most effectively in a multi-step process representing the transient characteristics seen in Figure 4.14. It is helpful to first recognize that the linear representation of the unstable pole in the *s*-plane moves further into the right half plane as the heat release is increased; this corresponds to an increase in self-excited loop gain. This indicates that more control authority/gain is required at higher heat release conditions to stabilize the system. The core of the hypothesis is that the operating point changes, causing the self-excited loop gain to decrease, when the flame is excited by oscillating pressure. This effect occurs as a result of unmodeled low frequency dynamics with long time constants and is explained in [Khanna, 2001]. If we begin with the uncontrolled system at a high heat release operating condition and refer to Figure 4.15, the following physical phenomena occur:

- Without control, the self-excited linear system pole resides in the right half plane at operating point T1. Recognize that the *operating point* may change while the *operating conditions* remain constant.
- When the feedback controller stabilizes the system, the controlled system pole moves to location 2, still at the T1 operating point. Because the flame is no longer oscillating, the operating point begins to change causing gain in the self excited loop to increase. The controlled system pole gradually moves to location 3.
- When the self excited loop gain reaches operating point T2, the gain applied by the feedback controller is no longer sufficient to stabilize the system, and the flame begins to oscillate as the limit cycle grows.
- After sufficient oscillation occurs, the operating point moves back to T1, the self-excited loop gain is reduced, and the feedback control gain can again stabilize the system. The process repeats indefinitely.

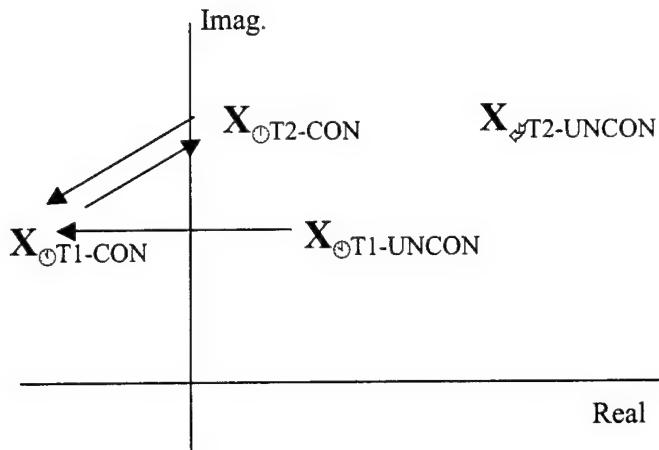


Figure 4.15 Controlled Pole Locations

It is clear that this occurs because there is insufficient gain in the feedback controller to keep the system stable after initial control raises the self-excited loop gain. Indefinite stabilization only occurs when the controller has enough gain to stabilize the self-excited loop when the flame is not oscillating. Position 4 represents the *uncontrolled* self-excited dynamics if the flame were not oscillating, at the original operating condition and the new operating point. Because a stable flame increases the loop gain of the self-excited system, the controller must be capable of stabilizing the pole as it exists in location 4. For the experimental Rijke tube examined here, this only occurs in the range of lower equivalence ratios, based solely on the power of the acoustic actuator.

To prove the dependence of loop gain on flame oscillation, a burner stabilized flame, similar to that in the Rijke tube, was excited at various probe amplitudes while the gain of the flame transfer function was recorded. The following table shows the results.

Probe Gain (V) @ 130 Hz	4.4.3.1.1	FRF Amplitude (dB) @ 170 Hz
0.0		0.86
0.05		0.82
0.1		0.58
0.2		-1.32
0.3		-3.00

The FRF amplitude of the heat release transfer function translates directly into self-excited loop gain. It is therefore clear that changes in acoustic pressure oscillations affect the gain of the self-excited loop. It should be noted that this effect can also be observed in the low equivalence ratio cases that were stabilized. For example, in the upper trace of Figure 4.13, near the zero second point, a change in the level of the controlled oscillation is observed. This is undoubtedly due to the changing of the heat release dynamics following the dramatic change in oscillation amplitude after control is first applied.

Because the actuator does not have the power to stabilize the system indefinitely, no controller design can provide stabilizing control. However, by simply reducing the gain of

the feedback controller, a new stable limit cycle, at a lower amplitude than the uncontrolled limit cycle, can be achieved without intermittency. For adaptive controllers, this can be implemented using an appropriately large leakage parameter so that the adaptive weights cannot reach their optimal solution. Another solution is to add a secondary probe signal to the control signal to eliminate the searching behavior. This intentionally oscillates the flame so that the self-excited loop gain is low enough that the controller can maintain stabilizing control at the limit cycle frequency. Each of these techniques strives to reduce the limit cycle amplitude as much as possible using all of the available actuator authority, without allowing intermittency.

4.5 Stability and Operating Constraints of Adaptive LMS-Based Feedback Control

Although the LMS-based feedback controllers discussed in the previous sections have been reliable in practice, the stability guarantees of the feedforward LMS implementations do not carry over to the feedback implementations. The following sections examine two distinct mechanisms for possible instability: feedback loop instabilities and algorithm divergence. Contrary to feedforward adaptive control, feedback loop instabilities can be created when adaptive controllers are used in structures containing feedback loops. This mechanism is examined as a function of plant estimation error for both stable and unstable systems. A second mode of instability can occur when the plant estimate is inaccurate, causing the filtered-X LMS algorithm to diverge. This section shows why the conventional interpretation of acceptable plant estimation errors is incorrect in a feedback setting and presents a method for computing the correct gradient filter.

We begin our investigation of robustness of LMS-based feedback algorithms by first considering the case of controlling stable systems subject to exogenous disturbances. For the disturbance rejection problem, adaptive feedforward control has received considerable attention over the last few decades, particularly with application to active noise control [18]. The popular filtered-X variant shown in Figure 4.16 has clear advantages over comparable feedback structures.

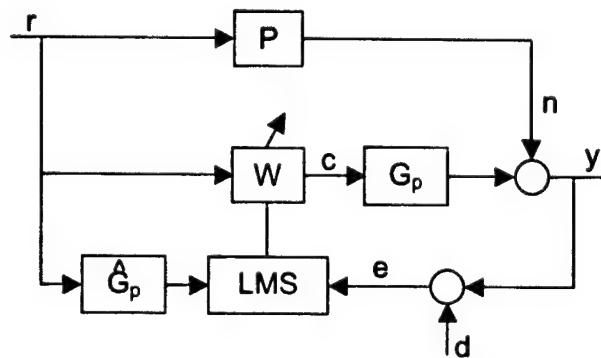


Figure 4.16 Adaptive Feedforward Control

Due to its feedforward architecture, the control system is inherently stable when W is an FIR filter and the adaptation is sufficiently slow. In addition, the filtered-X LMS

algorithm ensures convergence for physical systems whose adaptive filter output encounters a dynamic system before being sensed by the error signal.

The LMS algorithm updates the future weights of the adaptive filter in response to a scaled, instantaneous measurement of the local gradient (derivative of the squared error) as

$$W_{k+1} = W_k - \mu \nabla_k \quad (4.17)$$

The “filtered-X” moniker indicates that the reference signal must be filtered by an estimate of the control-to-error dynamics in order to accurately estimate the correct gradient at the error signal. This can be seen by computing the gradient directly as the derivative of the squared error with respect to the weight vector as shown below

$$\begin{aligned} E(z) &= -N(z) - G_p(z)W(z)R(z) \\ \nabla E(z) &= \frac{dE(z)}{dW(z)} = -G_p(z)R(z) \\ \nabla e_k &= -G_p * \vec{r}_k \\ \nabla_k &= \nabla e_k^2 = 2[\nabla e_k]e_k = -2(G_p * \vec{r}_k)e_k \end{aligned} \quad (4.18)$$

where * denotes the convolution operation in the sample domain.

It is clear that the reference input (r) must be filtered by the plant before being multiplied by the error to compute the instantaneous gradient. Since the plant itself is not available in physical applications, an estimate of that plant transfer function is typically used. It is well known [Elliott et. al., Kuo and Morgan, Morgan] that this estimate must be accurate to within 90° of phase when compared with the actual plant to ensure convergence of the algorithm. If the estimate is incorrect by more than 90° , the algorithm will search in the wrong direction and eventually diverge.

One of the primary limitations of practical implementation of the adaptive feedforward controller of Figure 4.16 is the requirement for an uncontrollable, coherent reference signal. The LMS algorithm assumes that r is highly correlated with n (the disturbance to be canceled) and that the output of the adaptive filter (c) does not influence r . If the former is violated the control performance suffers; if the latter is not satisfied, a feedback loop is introduced that might become unstable during adaptation. In active noise control applications, it is often difficult to obtain a reference signal that is both coherent with the disturbance and not influenced by the controller. Adaptive feedback control attempts to remedy this problem.

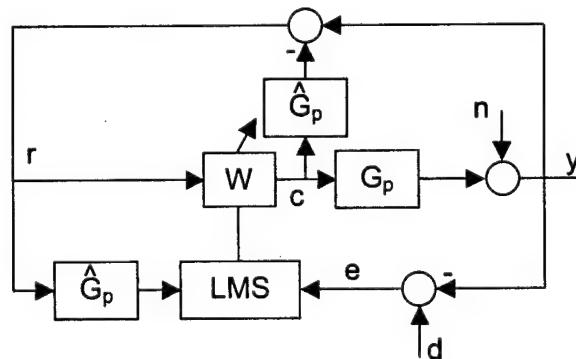


Figure 4.17 Filtered-E Control

Figure 4.17 illustrates a typical filtered-E controller arrangement [8], [14]. This arrangement employs the filtered-X LMS algorithm with an estimate of the actual control-to-error path \hat{G}_p . The external disturbance (n) enters at the error sensor along with the filtered control output as in Figure 4.16. Since there is not another reference signal available, the error is used to estimate the reference signal (r) by subtracting the influence of the controller from the error signal sensor.

$$r = (n + G_p c) - \hat{G}_p c \quad (4.19)$$

In view of (4.18), if $\hat{G}_p = G_p$ the reference signal is exactly equal to the disturbance to be canceled and is therefore coherent and uncontrollable. Significant difficulties can arise when \hat{G}_p deviates from G_p .

Before continuing, it should be noted that Figure 4.18 illustrates the same system of Figure 4.17, only redrawn. This reformulation of the filtered-E block diagram into a feedforward, filtered-X style system, clearly shows the inherent dependencies that the system has on the adaptive filter. Here we see that the adaptive filter is a part of the control-to-error path and can also adapt poles. The effects of these dependencies on the stability of filtered-E controllers are the focus of this section.

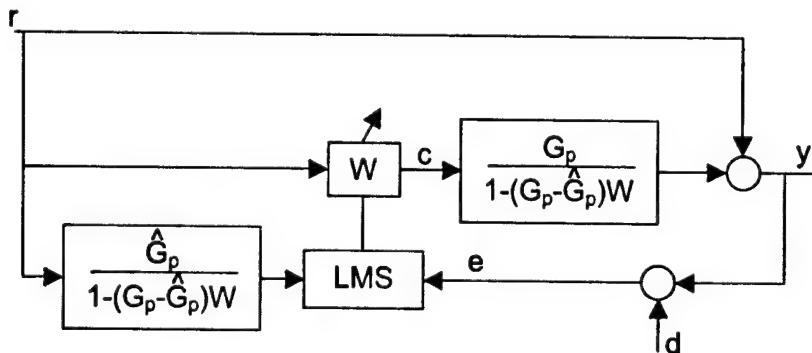


Figure 4.18 Filtered-E Control Redrawn

4.5.1 Feedback Loop Instabilities

Simply using an adaptive filter in a control system does not imply global stability. In fact, if it is employed in an arrangement as in Figure 4.17, stability can not be guaranteed during adaptation. Consider the system of Figure 4.17 at a moment in time when the adaptation is slow, or has stopped. A transfer function expressing the system input to output relationship is

$$\frac{y}{n} = \frac{1 + \hat{G}_p W}{1 + \hat{G}_p W - G_p W} \quad (4.20)$$

If \hat{G}_p is exactly equal to G_p the denominator vanishes, leaving a strictly feedforward system. Assuming W is an FIR filter and the control-to-error path is a stable system, the system dynamics are guaranteed to be stable regardless of W . If however, \hat{G}_p is not equal to

G_p at any frequency, the system poles can become unstable and are defined by the denominator of (4.20).

There are many classical control techniques that can be used to analyze the stability of linear systems including root locus, Routh Hurwitz, Nyquist plots, and Bode plots. For experimental systems where dynamics are not explicitly known, the most convenient and intuitive of these is the Bode analysis. The Bode stability criteria uses the open loop frequency response function to predict closed loop performance. The open loop transfer function of Figure 4.17 can be expressed as

$$\frac{G_p W}{1 + \hat{G}_p W} \quad (4.21)$$

At frequencies where the magnitude of (4.21) is greater than unity for phase values equal to 360 degree multiples (for positive feedback), the loop will be unstable. It should be noted that the Bode analysis assumes that the open loop transfer function is stable, i.e. that the roots of $(1 + G_p W)$ are in the left half plane. If they are not (determined by examining the open loop frequency response of $G_p W$), the Nyquist criteria provides a similar stability analysis for the closed loop system assuming the number of unstable roots is known explicitly.

It should be clear from (4.20) and (4.21) that any loop instability resulting from unstable poles of (4.20) is a function of both the adaptive filter and the error between the plant and its estimate. This is a system dependent phenomenon which cannot be predicted by assuming W only reaches a fraction of the optimal solution as suggested by [Leboucher et. al.]. The optimal adaptive filter (found when the error is set to zero) is

$$W_{opt} = -\frac{1}{\hat{G}_p} \quad (4.22)$$

Assuming a plant estimate error exists, and the adaptive filter has reached the optimal, (4.20) can still become unstable at a specific frequency even though the numerator is zero. Likewise, it is possible that the denominator of (4.20) is stable, and the system has driven the error to zero with no unstable roots. The stability of the system must be examined on a case by case basis. It will be a function of the frequency dependent error between the plant and plant estimate as well as the size and shape of the adaptive filter, regardless of whether it has reached the optimal solution.

Equation (4.22) illustrates that the optimal solution is only a function of the plant estimate, which we have obtained *a priori*. This raises important questions about the purpose of adapting if the optimal solution is already pre-determined. For broadband disturbance suppression, (4.22) must be satisfied at every frequency. Typically this is an impractical modeling task for non-minimum phase systems or finite duration FIR filters. Adaptation of W will result in the best compromise for control by inverting at the specific frequencies contributing most to the mean squared error. Tracking changes in those disturbance frequencies also presents a good argument in favor of adapting W which will ensure the optimal is maintained.

4.5.2 Algorithm Divergence

Conventional interpretation of the plant estimation error for the filtered-X LMS algorithm indicates a 90° error between the control-to-error path and it's estimate is

tolerable. This discussion illustrates that this interpretation is not valid for the filtered-E system of Figure 4.17. First consider that

$$r = y - \hat{G}_p c = n + (G_p - \hat{G}_p) c \quad (4.23)$$

follows from Figure 4.17 and (4.18).

Assuming slow adaptation, it is possible to rearrange Figure 4.17 such that the reference (r) is the system input as well as the exogenous disturbance; Figure 4.18 results. The feedforward structure of Figure 4.18 facilitates the understanding of the *actual* plant error analysis as compared to the *conventional* plant error.

The *conventional* plant estimate error of this system is the difference in the control-to-error path and the estimate of the control-to-error path used for the "filtered-X" portion of the algorithm. From Figure 4.18, this can be seen to be

$$\angle \left(\frac{G_p}{1 - (G_p - \hat{G}_p)W} \right) - \angle \left(\frac{\hat{G}_p}{1 - (G_p - \hat{G}_p)W} \right) < 90^\circ \quad (4.24)$$

This is expressed equivalently as

$$\angle G_p - \angle \hat{G}_p < 90^\circ \quad (4.25)$$

and is hereafter defined as the *conventional plant estimation error* (which is equivalent to the plant estimation error in view of (4.18)) because of its common usage in all adaptive controller arrangements. This is consistent with the conclusions reached in (4.18) and represents the common interpretation of the plant error that must be satisfied to ensure convergence for filtered-X LMS control.

The filter preceding the LMS algorithm (the second term in (4.24)) is the conventional reference signal filter when employing the filtered-X control strategy of Figure 4.17 and Figure 4.18. Conventional thinking dictates that as long as (4.25) is satisfied, convergence to the optimal solution (minimum MSE) will continue. The instance when \hat{G}_p is exactly equal to G_p satisfies this constraint because the feedback path is eliminated from the actual control-to-error path of Figure 4.18. However, when \hat{G}_p is not precisely equal to G_p , a feedback path is introduced that is a function of the adaptive filter W . We will now consider the case during adaptation, when \hat{G}_p has an arbitrarily small amount of estimation error with respect to G_p , causing the feedback loop to exist.

As before, the LMS algorithm updates the weights based on the negative gradient as in (4.17). For simplicity we examine the error vector gradient which defines the filter used to filter the reference signal in the filtered-X algorithm. The *actual* error gradient of the filtered-E system is now computed based on the nonzero feedback path in the physical control-to-error transfer function of Figure 4.18.

$$E(z) = D(z) - X(z) - \frac{G_p(z)W(z)X(z)}{1 - (G_p(z) - \hat{G}_p(z))W(z)}$$

$$\nabla E(z) = \frac{dE(z)}{dW(z)} = \frac{-G_p(z)}{1 - 2(G_p(z) - \hat{G}_p(z))W(z) + (G_p(z) - \hat{G}_p(z))^2 W(z)^2} X(z) \quad (4.26)$$

The filter used to filter the reference signal in (4.26) must now be compared to what is conventionally used by the filtered-x in Figure 4.18 to estimate the gradient (without the explicit dependence on z) as

$$\angle \left(\frac{G_p}{1 - 2(G_p - \hat{G}_p)W + (G_p - \hat{G}_p)^2 W^2} \right) - \angle \left(\frac{\hat{G}_p}{1 - (G_p - \hat{G}_p)W} \right) < 90^\circ \quad (4.27)$$

Equation 4.27 represents the difference in the filter that should be used to filter the reference based on the actual gradient estimate, and the filter that is conventionally used in the filtered-X formulation of Figure 4.17.

The well-known 90° phase limitation between the physical control-to-error path and the plant estimate has historically been derived only for a single sinusoid [Morgan, Snyder and Hanson, Elliot et. al.]. Assumptions made for these derivations rely on a filtered-X problem formulation where there is no dependence of the control-to-error path dynamics on the adaptive filter. An examination of Figure 4.18 shows that unlike the feedforward situation, where we adapt an FIR filter $W(z)$ that linearly affects the control signal, the filtered-E structure is essentially adapting an IIR filter $H(z)$ given by

$$H(z) = \frac{G_p(z)W(z)}{1 - (G_p(z) - \hat{G}_p(z))W(z)} \quad (4.28)$$

where the parameters to be adapted are contained in the FIR transfer function $W(z)$. To minimize the mean square error, the gradient is still computed as in (4.26) and can be expressed as the partial derivative of $H(z)$ with respect to $W(z)$

$$\frac{\partial H(z)}{\partial W(z)} = \frac{G_p(z)}{1 - 2(G_p(z) - \hat{G}_p(z))W(z) + (G_p(z) - \hat{G}_p(z))^2 W(z)^2} \quad (4.29)$$

Clearly, this is much different than the transfer function used in the filtered-E algorithm, which is the lower left block of Figure 4.18, and shown as the second term in (4.27). The question then arises as to how much can the gradient used by the filtered-E algorithm differ from the true gradient and yet still produce convergence to the optimal? It would seem to be extremely difficult to answer this question globally, since the equations describing nominal trajectories in the weight space are nonlinear. Certainly near the optimum, W_{opt} , the cost will be quadratic and we can approximate $H(z)$ in a neighborhood of the optimum as a transfer function that is affine in the parameters W as a Taylor's Series expansion about the optimal:

$$H(z) = H(z) \Big|_{W_{opt}} + \frac{\partial H(z)}{\partial W(z)} \Big|_{W_{opt}} (W - W_{opt}) \quad (4.30)$$

The first and last terms of (4.30) are moved to the forward (disturbance) path between the reference and error sensor in Figure 4.18, while the control path containing W is altered by the linearized control-to-error dynamics of $\partial H(z) / \partial W(z) \Big|_{W_{opt}}$. Because this filter is not a function of the changing adaptive filter, the standard 90° phase error analysis for the filtered-X algorithm now applies. The conclusion is that as long as the phase of the gradient filter does not depart from the phase of $\partial H(z) / \partial W(z) \Big|_{W_{opt}}$ by more than 90° at the frequency of interest, the algorithm will converge to the optimal solution. If the phase difference is

greater than 90° , then the optimal cannot be approached, and the algorithm will diverge from the optimal. Since we have no idea *a priori* as to the value of W_{opt} , it makes sense to use (4.29) as the gradient filter. It should be noted that in our simulations, we witnessed divergence over a wide range of the weight space whenever the gradient filter used by the filtered-E algorithm differed from Equation (4.29) by more than 90° . Note that the negative sign computed in (4.26) cancels with the negative sign of (4.17) when the FX-LMS algorithm is implemented so that the sign of the filter agrees with the physical control-to-error path predicted by (4.30) from (4.29).

Returning to the conventional implementation of the filtered-E controller, the inequality expressed by (4.27) defines the difference in the actual plant estimation filter and the conventional plant estimation filter that is typically used. It should be clear that when comparing (4.27) and (4.25), the errors will result in different predictions of algorithmic stability as a function of conventional plant error. It is therefore conceivable that a small plant estimation error in *conventional* terms that satisfies (4.25), could produce a large error that violates the *actual* estimation constraint (4.27), causing an unexpected algorithm divergence, but one that is predicted by (4.27).

Given this result, one might assume that since the actual gradient filter is known, it can be used in place of the conventional filter to ensure that the plant estimate is accurate enough to satisfy (4.27). Since the actual plant estimation filter of (4.26) is a function of the difference in the actual plant dynamics (G_p) and the estimate of the dynamics (\hat{G}_p), this is impossible. If the actual dynamics were known exactly, they could be used to eliminate the feedback path resulting in a strictly feedforward system. The tacit assumption in conventional adaptive feedback control is that the estimate exactly equals the actual plant and there is no error, thereby turning (4.27) and (4.20) into (4.25) and the numerator of (4.20). In practice we know there exists some finite error between the estimate and the actual control-to-error path. However, this can only be analyzed in a simulation where the error can be explicitly controlled and the actual plant is known exactly.

It is interesting to examine the behavior of (4.27) as a function of the adaptive filter W . Evaluating (4.27) when $W=0$, results in

$$\angle \left(\frac{G_p}{1 - 2(G_p - \hat{G}_p)W + (G_p - \hat{G}_p)^2 W^2} \right) - \angle \left(\frac{\hat{G}_p}{1 - (G_p - \hat{G}_p)W} \right) \Big|_{W=0} = \angle G_p - \angle \hat{G}_p < 90^\circ \quad (4.31)$$

Also noteworthy is the result when (4.27) is evaluated when the adaptive filter reaches the optimal solution of (4.22).

$$\angle \left(\frac{G_p}{1 - 2(G_p - \hat{G}_p)W + (G_p - \hat{G}_p)^2 W^2} \right) - \angle \left(\frac{\hat{G}_p}{1 - (G_p - \hat{G}_p)W} \right) \Big|_{W=-\frac{1}{\hat{G}_p}} = 0 \quad (4.32)$$

Therefore at the inception of control, assuming a zero initial condition on the adaptive filter, the conventional gradient error estimation is valid. In addition, when the adaptive filter reaches the optimal solution, there is no estimation error, regardless of the choice of \hat{G}_p ! As a result, divergence of the adaptive algorithm due to inaccurate plant estimation should only be expected *during* adaptation. Although, practically speaking, a finite length FIR filter will not likely have the ability to invert the control-to-error path estimate at every frequency in the controllable bandwidth. Therefore gradient divergence as predicted

by (4.27) will always remain a practical concern when implementing adaptive feedback controllers like the filtered-E.

4.6 Practical Applications and Considerations

4.6.1 Interrelated Instabilities

The two instability modes presented above are connected to each other after an initial transient. Each instability can initiate independently, but because of the inherent dependence on W in the gradient estimate, they will eventually influence each other. Suppose that a conventional control-to-error path estimation error exists, but that (4.27) is satisfied and the adaptive filter is converging toward the optimal solution. Because an error exists, the poles of (4.20) exist and are moving as a function of the error and the adaptive filter. It is possible for one of these poles to leave the unit circle during adaptation and create a feedback loop instability. This can happen independent of violating (4.27). However, as the loop instability grows in magnitude, the mean squared error will increase and cause the adaptive filter to respond. In view of (4.27), a changing adaptive filter will affect the actual gradient filter error, and potentially cause it to exceed 90° . Therefore a feedback loop instability can cause the phase error to exceed the 90° criterion established by (4.27) resulting in algorithm divergence. Alternatively, the adaptation could potentially re-stabilize the system! It is also important to note that because of the dependence of the control-to-error path on the adaptive filter, it is possible that the phase estimation error defined by (4.27) could be self-correcting during adaptation. In other words initial error in excess of 90° may not guarantee algorithm divergence if the adaptive filter changes the control-to-error path in such a manner that the actual phase error is reduced.

Although each instability mechanism can initiate independent of the other, they are ultimately codependent through the adaptive filter magnitude. That is, if one grows without bound, the other will follow. For this reason, it is virtually impossible to ascertain the mechanism of instability in experimental applications because the time scales are typically too fast. The simulations to follow will illustrate these phenomena in closer detail.

4.6.2 Filtered-U LMS Algorithm

The filtered-E controller examined above can be viewed as a special case of the more generalized filtered-U LMS algorithm [Wang and Ren]. The filtered-U algorithm, often employed in duct noise control problems, is illustrated in Figure 4.19. Here there are two control sensors: the "upstream" reference microphone detects the original source disturbance (v) and generates the input to the algorithm (x) while the downstream error microphone senses the noise (n) after being altered by the duct acoustics existing between those two sensors, represented by P . Because the two sensors are in the same duct, when the control signal (c) is applied to the system it influences the error signal through the control-to-error path dynamics (G_p) as well as the upstream reference sensor through a feedback path, F .

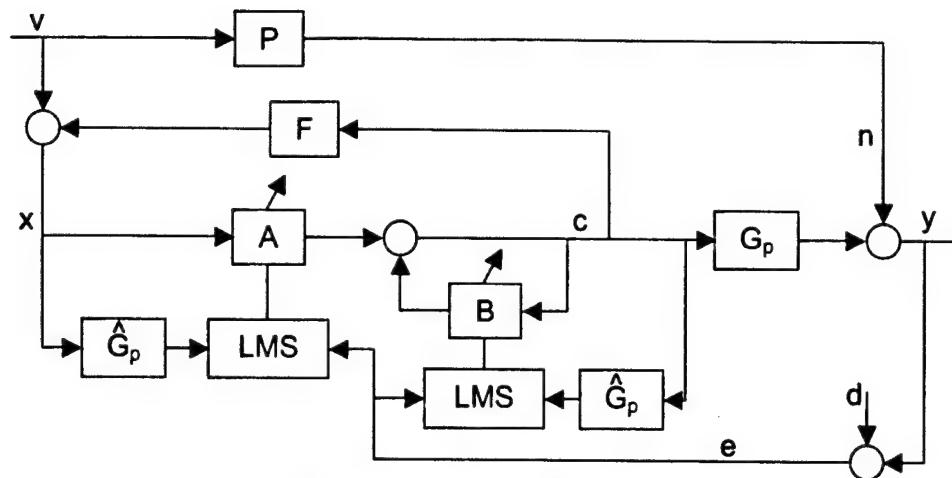


Figure 4.19 Filtered-U Control

To see the similarities between the filtered-U and filtered-E structures we need only to consider the case when the reference and error signals are identical. When this occurs, physically, there is only one microphone in the duct and therefore no reference signal "v". The transfer function F disappears, P becomes unity, and the reference signal, x , is the same as the error microphone measurement, y . The disturbance now appears as an exogenous input to the error sensor. These changes and the resulting system are depicted in Figure 4.20.

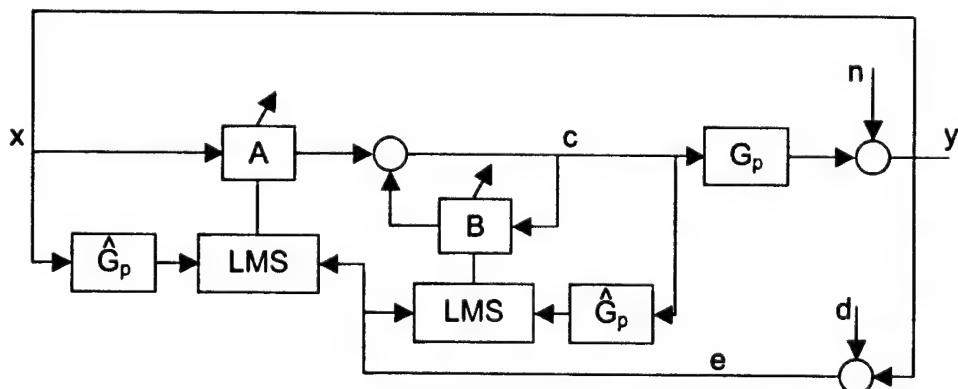


Figure 4.20 Filtered-E Adaptive IIR Control

Comparing Figure 4.17 and Figure 4.20, an obvious difference exists in the controller structure. Figure 4.17 (the filtered-E) uses a fixed estimate of the control-to-error dynamics to compensate for the feedback path created by using the error signal as the reference. Figure 4.20 illustrates a filtered-E controller structure as well, where the feedback path is compensated adaptively by the adaptive IIR form. Equating the transfer functions of Figure 4.17 and Figure 4.20, we see that the two systems are identical when $A = W$ and $B = -WG_p$. However, because of the adaptive IIR structure there is no guarantee that these will be the optimal filters or that they will ever be reached [Kuo and Morgan].

As with the adaptive feedback analysis, the filtered-U system of Figure 4.19 has the same loop stability limitations. Because the reference signal is controllable and the adaptive feedback component does not begin at the optimal solution, the possibility for feedback instabilities exists. In addition, because of the feedback to the reference signal and the existence of an adaptive feedback filter (B), the conventional gradient estimate cannot be accurate during adaptation as shown next.

Block diagram algebra of Figure 4.19 reveals:

$$E(z) = D(z) - P(z)V(z) - \frac{G_p(z)A(z)}{1 - B(z) - A(z)F(z)}V(z) \quad (4.33)$$

$$V(z) = X(z) - F(z)C(z)$$

Examining the actual gradient of the error signal (which is proportional to the gradient of the cost function) with respect to the forward path filter (A) results in

$$\frac{dE}{dA} = \frac{BG_p - G_p}{(1 - B - AF)^2}V = \frac{-G_p(1 - B)}{(1 - B - AF)^2}X + \frac{FG_p(1 - B)}{(1 - B - AF)^2}C \quad (4.34)$$

Conventional application of the filtered-U algorithm uses an estimate of the control-to-error dynamics (\hat{G}_p) to filter the input of the adaptive filter for updating the weights. In the case of the forward component of the IIR filter employed here, the input is x and it would typically be filtered by \hat{G}_p . In fact, the actual gradient estimation includes the feedback path and is also a function of the output signal c as shown in (4.34). A similar analysis can be performed for the partial derivative of the error signal with respect to the feedback adaptive filter B, yielding similar results. It is therefore clear from (4.34) that the conventional and actual gradient estimation errors can differ quite significantly, and may unpredictably result in algorithm divergence even when (4.25) is satisfied.

4.6.3 Online System Identification

Online system identification is typically employed to track changes in the control-to-error dynamics over time. For adaptive feedforward systems with time varying control-to-error transfer functions, this technique may be required to maintain the correct gradient estimate in the LMS algorithm. Figure 4.21 [Kuo and Morgan] illustrates one possible technique for performing the online system identification for a filtered-X feedforward system. It is assumed that the plant G_p is time varying and must be estimated by a continually updated \hat{G}_p .

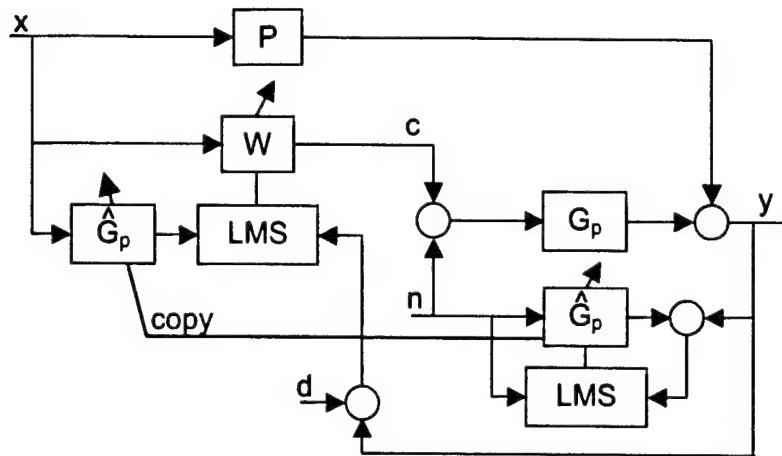


Figure 4.21 Online System Identification

As shown above for filtered-E control, the actual control-to-error path is a function of the adaptive filter, and is therefore changing with time. The transfer function coefficient of X in (4.26) represents the actual control-to-error path estimate that should be used to filter the input to the LMS algorithm. The question then arises, "if the actual gradient equation is known, why can't it be used in the filtered-X update algorithm". Upon inspection of (4.26), it should be clear that G_p is never known exactly. \hat{G}_p is used to estimate G_p ; and if the assumption is that there is no estimation error $\hat{G}_p = G_p$ and (4.27) becomes (4.25). However, in practice there is always an estimation error, no matter how small, that makes (4.27) the correct gradient error criteria. Therefore, without knowing G_p exactly, the $(G_p - \hat{G}_p)$ terms of (4.26) cannot be computed directly.

Another alternative for generating a more accurate estimate of the actual gradient is the online identification of the gradient. Once the loop is closed (the adaptive filter becomes nonzero), an online identification procedure will take into account the feedback loop and the adaptive filter. However, system identification of any input/output relationship of Figure 4.17, will never result in the required transfer function coefficient in (4.26). This is shown clearly by considering that the characteristic equation of Figure 4.17 is the denominator of (4.20). Therefore, schemes directly using the results of online system identification cannot result in the transfer function of the first term in (4.27), which has a quadratic denominator.

A similar argument can be made for an online system identification employed in the filtered-U system of Figure 4.19. The quadratic denominators of (4.34) prevent an online system identification procedure from ever yielding the proper frequency response. [Kuo and Morgan] and [Feintuch] proposed that the feedback terms resulting from F and B in (4.34) were negligible. In this case the gradient estimates for both adaptive filters become $\hat{G}_p = G_p$; a result that ignores the recursion of both B and F . However, an online system identification procedure cannot avoid identifying the physical feedback paths, F and B , in Figure 4.19. Therefore the frequency response resulting from the online identification process (Figure 4.21) of the filtered-U algorithm (Figure 4.19) is

$$\frac{G_p}{(1 - B - AF)} \quad (4.35)$$

which is not an accurate representation of their proposed gradient filter *approximation*, G_p .

4.6.4 Extension to Self-Excited Systems

Our previous analysis has focused on systems having exogenous disturbances acting on stable plants. We now consider how this analysis carries over to the main problem of controlling thermoacoustic instabilities. Figure 4.22 illustrates the adaptive feedback controller applied to a self-excited instability in the form of a thermoacoustic instability. G_a represents the acoustic frequency response of the combustor while G_f represents the nonlinear flame dynamics. G_{act} represents the acoustic transfer function plus actuator dynamics between the control signal to error signal, typically including time delay. Unfortunately, the same instability mechanisms for the stable plant control exist for the unstable plant control. The characteristic equation for the adaptive feedback control system of Figure 4.22 is represented by

$$1 - G_a G_f - W(G_{act} + \hat{G}_p - G_a G_f \hat{G}_p) \quad (4.36)$$

There are no guarantees on the stability of the roots of (4.36). The adaptive filter could easily cause a root of this equation to become unstable during adaptation, regardless of the choice of \hat{G}_p . This is unlike the stable plant control where the feedback loop is canceled if $\hat{G}_p = G_p$. Without accurate knowledge of the self-excited system dynamics, it is impossible to limit the roots of (4.36) to strictly stable values. It should be noted that when $W = W_{opt}$ the roots of (4.36) are the roots of the denominator of G_{act} . But no guarantees can be made on reaching the optimal solution, especially if the gradient estimate is incorrect.

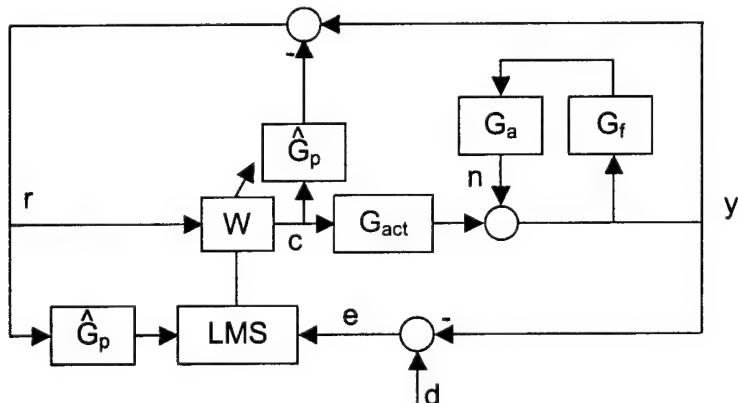


Figure 4.22 Self-Excited System Filtered-E Control

The gradient of the cost function is even more complex for the self-excited case. The filter to which the input to the LMS algorithm should be applied will be a function of the adaptive filter, plant estimate and actual plant as before, but will also be influenced by the self-excited system dynamics. Computation of the actual gradient of the filtered-E self-excited system should not be required in order to recognize that the conventional estimate of G_{act} will be inaccurate with respect to the actual gradient.

4.7 Simulation

Two simulations have been designed that illustrate the two specific instability mechanisms independently. Each simulates the stable disturbance rejection system of Figure 4.17 using different plants and plant estimates. As set forth in the previous discussions, the *conventional* plant error refers to the criterion of (4.25) which was shown to be inaccurate for

adaptive feedback controllers. Equation (4.27) represents the *actual* estimation error criterion.

4.7.1 Feedback Loop Instability

This simulation illustrates a case where the conventional plant error is in excess of 90° at many frequencies, but no algorithm divergence is observed. Instead, an unstable feedback loop is generated due to a pole of (4.20) leaving the unit circle.

The plant shown in Figure 4.17 was chosen to have unity magnitude and a linear delay of 25 samples while the plant estimate was chosen to have unity magnitude and 19 samples of delay at a sample rate of 2000 Hz. The exogenous disturbance was chosen to be a single sinusoid at 32 Hz with additive white noise at a lower level at every other frequency. The adaptive filter was a 2 weight FIR filter with a convergence parameter of 0.00003.

The difference in phase between the plant and estimate increased almost linearly, reaching an excess of 1000° of phase error by 1000 Hz, thus allowing for the possibility of divergence of the weights due to a conventionally inaccurate system identification. Figure 4.23 illustrates the uncontrolled (solid) and controlled (dotted) power spectra of the tonal disturbance from the stable plant at 32 Hz. Initially, the tone is suppressed with only two adaptive filter weights but the loop gain that accompanies the optimal adaptive filter causes a loop instability at 810 Hz. This is accurately predicted by the Bode gain/phase relationship in that the open loop frequency response magnitude is in excess of 0 dB at the 810 Hz phase crossover frequency.

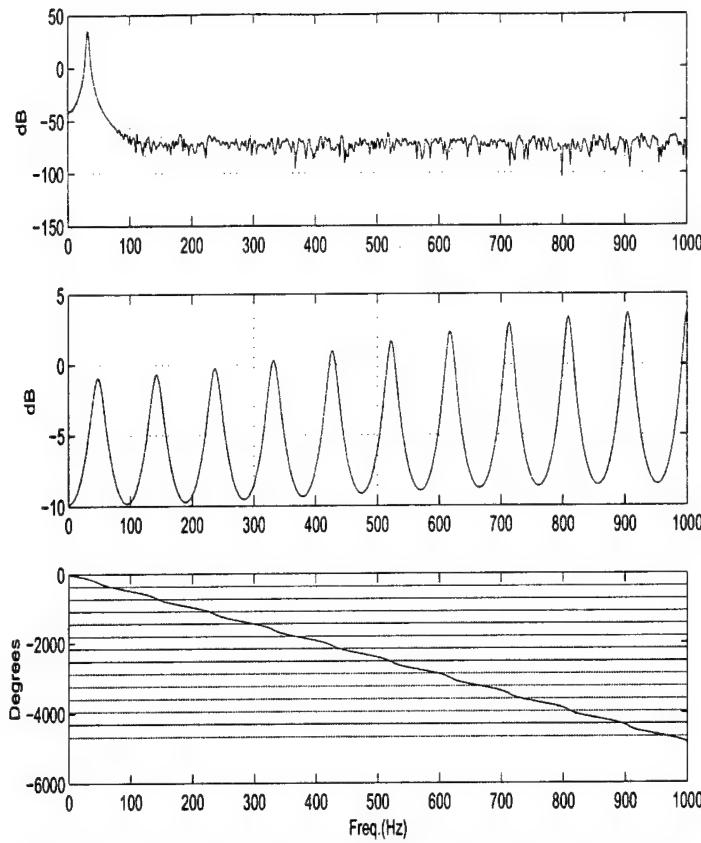


Figure 4.23 Feedback Loop Instability and Bode Plot Prediction

This adaptive feedback controller diverges because of the feedback loop instability, not divergence of the algorithm. Increasing the number of adaptive filter weights reduces the gain of the adaptive filter at 810 Hz thus eliminating that instability. However, the added phase and filter complexity causes poles at other frequencies, away from the disturbance, to become unstable before convergence is achieved. For this simulation, the *conventional* phase error between the plant and estimate is still in excess of 90° at many frequencies throughout the control bandwidth but the algorithm *never* diverges because the actual phase error is less than 90° .

4.7.2 Algorithm Divergence

When the adaptive filter is at zero or its optimal solution, the *actual* plant error is either equivalent to the *conventional* error or is degenerate. Therefore we are only concerned with the times during adaptation when neither of these conditions are satisfied. In addition, we are interested in illustrating a case where the *conventional* plant estimation error is arbitrarily small but the *actual* plant estimation error exceeds the constraints of (4.27). Finally, it is important to continue to differentiate the gradient based algorithm divergence from the loop instability presented above. This simulation accomplishes each of these goals.

Because the actual plant estimation error (4.27) established herein is a function of the plant estimation error as well as the adaptive filter, it is impossible to generalize the expected gradient error. For this particular simulation the plant estimate was chosen to be unity so

that the optimal adaptive filter was -1 . The control-to-error path was chosen to be a complex zero at 200 Hz with a damping ratio of .025 combined with a complex pole at 300 Hz with a damping ratio of .016. Due to the way the simulation was designed an additional delay was imparted to both the control-to-error path and its estimate (with a sample frequency of 2000Hz). These choices resulted in a *conventional* phase difference in the plant and plant estimate that was less than 4° at the disturbance frequency of 150 Hz. Therefore in a conventional interpretation where the tolerable error satisfies (4.25), there is no chance for gradient divergence at 150 Hz.

In a user controlled simulation environment, we have access to both the actual plant and the plant estimate. Therefore it is possible to compute the phase difference described by (4.27) directly for a variety of adaptive filters. In order to effectively visualize the weight space, a two weight adaptive filter was chosen. This is also typically sufficient to control a single tone disturbance. Figure 4.24 illustrates the *actual* gradient error as computed by (4.27), as a function of the two adaptive weights, at 150 Hz. The *'s represent the weight combinations that result in a gradient error of greater than 90° at 150 Hz; the absence of *'s represent areas where the *actual* estimation error is less than 90° . Recall that the *conventional* gradient error for the entire weight space, at 150 Hz, is less than 4° . The optimal adaptive solution of (4.22) is shown as a star at coordinates $(0.98, -1.77)$. Note that the *actual* gradient error at $(0,0)$ and the optimal, is within the 90 degree specification.

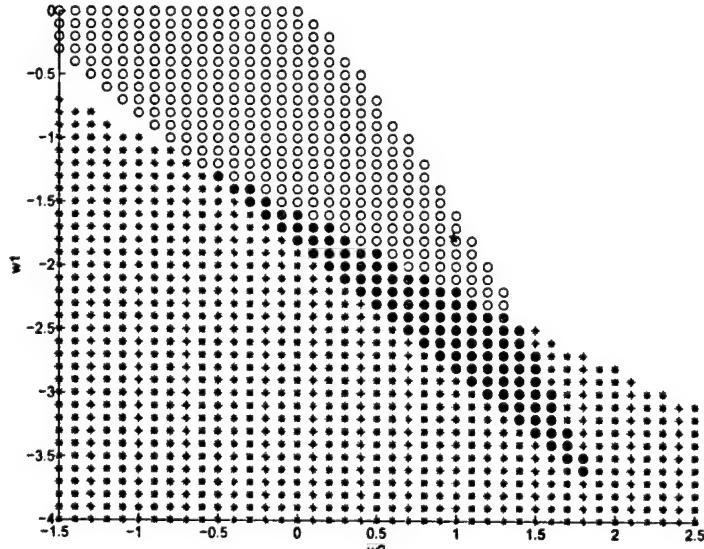


Figure 4.24 Gradient and Stability in the Weight Space

Figure 4.24 also illustrates loop stability. For every weight combination shown, (4.20) was evaluated for unstable roots. The O's represent areas in the weight space where the loop has no unstable roots; areas without O's have at least 1 unstable root. By inspection, it is possible to see a union where the feedback loop is stable but the actual gradient is greater than 90° at 150 Hz ((4.27) is not satisfied) while the *conventional* gradient is less than 4° ((4.25) is satisfied). If we choose an initial condition of the weights in this region (-2.8 and 1.3 for example), the algorithm diverges at 150 Hz but the feedback loops remain stable for some time. The upper portion of Figure 4.25 illustrates the error signal with time while the lower

portion shows the increase in amplitude of the 150 Hz tone after control has been applied. This example illustrates that the conventional interpretation of the plant phase error of (4.25) is insufficient to ensure convergence to the optimal solution during adaptation. Equation (4.27) accurately predicts the *actual* phase error that can be expected during adaptation.

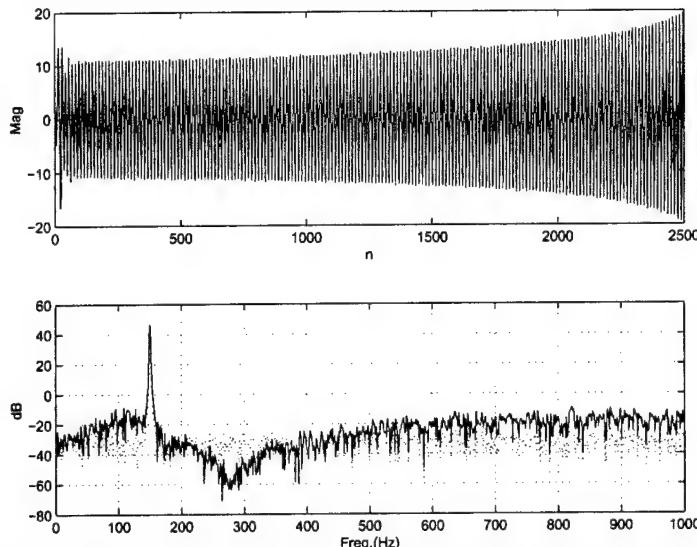


Figure 4.25 Gradient Divergence for 4 degrees of Plant Estimation Error

4.8 References

- [1] Annaswamy, A.M., O.M. El Rifai, M. Fleifil, J.P. Hathout, and A.F. Ghoneim, "A Model-based Self-tuning Controller for Thermoacoustic Instability," *Combustion Science and Technology*, **135**, 1998, pp. 213-239.
- [2] Billoud, G., M.A. Galland, C.H. Huu, and S. Candel, 'Adaptive Active Control of Combustion Instabilities,' *Combust. Sci. and Tech.*, **81**, 1992, pp.257-283.
- [3] Elliott S.J and P. A Nelson, "Active Noise Control," IEEE Signal Processing Magazine, Oct.1994, pp 12-35.
- [4] Elliott S.J., I.M. Stothers, and P.A. Nelson, 'A Multiple Error LMS Algorithm and Its Application to the Active Control of Sound and Vibration,' *IEEE Transactions on Acoust., Speech, and Sig.Proc.*, Vol.ASSP-35, No.10, Oct., 1987.
- [5] Eriksson, L.J. and M.C. Allie, "Correlated Active Attenuation System with Error and Correction Signal Input," U.S. Patent 5,206,911, Apr. 27, 1993.
- [6] Eriksson, L.J., M.C. Allie, and R.A. Greiner, "The Selection and Application of an IIR Adaptive Filter in Active Sound Attenuation," *IEEE Trans. Acoust. Speech, Signal Processing ASSP-35*, Apr. 1987, pp. 433-437.
- [7] Elliott S.J., I.M. Stothers, and P.A. Nelson, 'A Multiple Error LMS Algorithm and Its Application to the Active Control of Sound and Vibration,' *IEEE Transactions on Acoust., Speech, and Sig.Proc.*, Vol.ASSP-35, No.10, Oct., 1987.

[8] Eriksson, L.J., "Recursive algorithms for active noise control," *Proc. Int. Symp. Active Control of Sound Vib.*, 1991, pp. 137-146.

[9] Feintuch, P.L., "An adaptive recursive LMS filter," *Proc. IEEE*, 64, 1976, pp. 1622-1624.

[10] Gutmark, E., T.P. Parr, K.J. Wilson, D.M. Hanson-Parr, and K.C. Schadow, 'Closed-loop Control in a Flame and a Dump Combustor,' *IEEE Control Systems*, 13, 1993, pp.73-78.

[11] Kemal, A. and C. Bowman, 'Active Adaptive Control of Combustion,' *Proceedings of the IEEE Conference on Control Applications*, 1995, pp.667-672.

[12] Khanna, V. 'Dynamic Analysis of Unsteady Heat Release During Gaseous Combustion,' Ph.D. Dissertation, VPI&SU, August, 2001.

[13] Koshigoe, S., T. Komatsuzaki, and V. Yang, 'Adaptive Control of Combustion Instability with On-Line System Identification,' *Journal of Propulsion and Power*, 15, 1999, pp.383-389.

[14] Kuo, S.M. and D.R. Morgan, 'Active Noise Control Systems, Algorithms and DSP Implementations,' John Wiley and Sons, Inc., New York, NY, 1996.

[15] Leboucher, E., P. Micheau, and A. Berry, "An Index to Predict the Stability of Decentralized Adaptive Feedback Active Noise Control System," *Canadian Acoustics*, 28, Part 3, 2000, pp. 60-61.

[16] Mohanraj, R. and B.T. Zinn, "Numerical Study of the Performance of Active Control Systems for Combustion Instabilities," *AIAA 98-0356*, 36th Aerospace Sciences Meeting & Exhibit, Reno, NV, Jan., 1998.

[17] Morgan, D., "An Analysis of Multiple Correlation Cancellation Loops with a Filter in the Auxiliary Path," *IEEE Trans. Acoust. Speech, Signal Processing*, ASSP-28, 1980, pp. 454-467.

[18] Nelson, P.A. and S. J. Elliot, *Active Control of Sound*, Academic Press Ltd., London, England, 1992.

[19] Saunders, W.R., M.A. Vaudrey, and B. A. Eisenhower, U. Vandsburger, and C.A. Fannin, 'Perspectives on Linear Compensator Designs for Active Combustion Control,' *AIAA 99-0717*, 37th AIAA Aerospace Sciences Meeting and Exhibit, Reno, NV, Jan. 1999-2.

[20] Snyder, S.D. and C.H. Hansen, "The effect of transfer function estimation errors on the filtered-X LMS algorithm," *IEEE Transactions on Signal Processing*, 42, 1994, pp. 950-953.

[21] Vaudrey, M.A., W.T. Baumann, and W.R. Saunders, "Applying Adaptive LMS to Feedback Control of Thermoacoustic Instabilities," Submitted to *IEEE Transactions on Control Systems Technology*, June 2001.

[22] Vaudrey, M.A., W.R. Saunders, and B. A. Eisenhower, 'A Test-Based Methodology for Apriori Selection of Gain/Phase relationships in Proportional, Phase-Shifting Control of Combustion Instabilities,' *Proceedings of ASME Turbo Expo 2000*, Munich, Germany, May 2000.

[23] Vaudrey, M.A., W. T. Baumann, and W. R. Saunders, 'Stability and Operating Constraints of Adaptive LMS-Based Feedback Control', submitted to IEEE Transactions on Applied Control, May, 2001.

[24] Wang, A.K. and W. Ren, 'Convergence Analysis of the Filtered-U Algorithm for Active Noise Control,' *Signal Processing Journal* 73 (1999), pp.255-266.

Widrow, B. and S.D. Stearns, 'Adaptive Signal Processing,' Prentice-Hall, Inc., Englewood, N.J., 1985.

5 Pulsed Control

5.1 Introduction

Most control schemes for the suppression of thermoacoustic instabilities involve fuel modulation actuators of the on-off type. In the past, these actuators were primarily forced at the instability frequency [1] but recently, in an effort to reduce the cycle fatigue and required bandwidth of the actuators, numerous researchers have considered using subharmonic forcing [2,3]. The main objective of this chapter is to determine whether subharmonic control is effective due to nonlinear behavior in the thermo-acoustic system or can be explained by linear analysis. To accomplish this objective, it is necessary to explore the ramifications of using pulsed actuation, as opposed to linear proportional actuation, in the control loop. Although the effects on the flame of using pulsed control signals have been examined [4,5], the analysis of the effect of pulsing on the controlled system has not been investigated.

We show that from the standpoint of stabilizing the system, only linear analysis is required and the quantity that determines control effectiveness is the Fourier component of the actuation signal at the instability frequency. The major ramification of this fact is that actuators used for control must have a bandwidth that extends to the instability frequency, even when pulsed subharmonically. In addition, for the usual case of fixed-amplitude pulsing it is shown that true stabilization does not occur, but a new, much smaller limit cycle replaces the original limit cycle. The amplitude of this new limit cycle can be predicted using linear control theory.

The basic block diagram for the thermo-acoustic interaction that can cause large pressure oscillations in continuous combustors is shown in Figure 5.1. The acoustics of the combustor interact with the nonlinear heat release dynamics in a feedback arrangement that is generated by one of several physical instability mechanisms[6]. Under certain operating conditions, the system becomes unstable and an acoustic oscillation grows until limited by nonlinear effects, resulting in a stable limit cycle. The goal of the control loop is to stabilize the system, eliminating the oscillation. In the simplest configuration, the measured acoustic pressure is phase shifted and fed to an acoustic or fuel actuator. The idea of subharmonic control is to divide the frequency of the phase-shifted acoustic signal by an integer M , producing a $1/M$ subharmonic, which is then fed to an actuator. For example, $M = 2$ produces a one-half subharmonic, and will have a fundamental frequency of $1/2$ of the instability frequency. In practice, the subharmonic generation is most easily accomplished by counting zero crossings. At every M^{th} positive-going zero crossing, a trigger pulse is generated. This pulse is used to trigger a one-shot, which produces a pulse of specified time duration after every trigger pulse. Thus, the input to the actuator is a pulse train at the subharmonic period with a specified on-time. Of course, the preceding description also applies to forcing at the input frequency by setting $M=1$. Proportional fuel injectors are possible [7], but can be difficult and more costly to use than traditional on-off fuel injectors. Therefore it is necessary to examine both proportional and fixed height control signals.

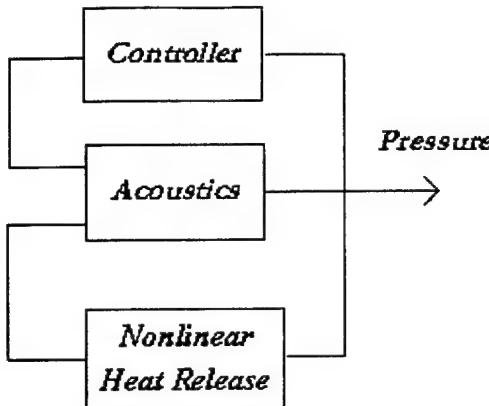


Figure 5.1 - Combustion Block Diagram

5.2 Analysis

5.2.1 Describing Functions

The simplest version of the controller we want to analyze is shown in Figure 5.2. ZCD is a positive-going zero crossing detector that produces a trigger pulse to a divide-by-M block which in turn produces a trigger pulse to a one-shot circuit that outputs a single pulse of duration w seconds. In practice, the zero-crossing detector will need a hysteresis characteristic to prevent accidental triggers due to noise. The basic controller is nonlinear and does not have a useful linearization. No matter how small the input signal, the output will not behave in a linear manner. However, we are interested in systems whose response is dominated by a single sinusoidal signal, a limit cycle. These instabilities normally occur at acoustic resonances and frequencies output by the controller at other than the instability frequency will be significantly attenuated by the acoustic plant, as they will not correspond to resonances. Thus, describing function analysis [8], which has traditionally been used to analyze limit cycling systems, is a natural candidate for a basic analysis of this control system. In this type of analysis, the nonlinear controller is replaced by a linear gain and phase at the frequency of excitation. If the input to the controller is $A \sin(\omega t)$ and the output component at frequency ω is $F(A, \omega) \sin(\omega t + \varphi(A, \omega))$, then the describing function is the complex gain given by

$$DF = \left| \frac{F(A, \omega)}{A} \right| \angle \varphi(A, \omega)$$

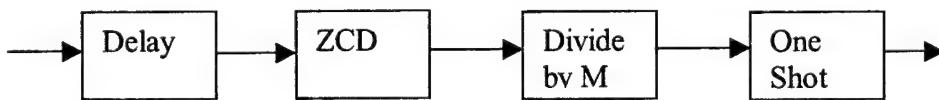


Figure 5.2 - Pulse Generator Block Diagram

When analyzing the system depicted in Figure 5.1, it must be noted that very little is known about the nonlinear heat release dynamics, even for simple flames. Even the order of the linearized model is unknown at present. To proceed, we make the assumption that at small signal levels the response of the heat release subsystem is dominated by a linear model

and linear analysis is valid. The experiments that follow are necessary to insure that the conclusions following from this assumption are observed in practice.

Under this assumption, single-frequency describing function analysis makes sense for the following reason. If the controller is effective, signal levels will be suppressed and the thermoacoustic system connected to the controller behaves linearly. Since superposition holds in the linear regime, an unstable oscillation will be unaffected by a control signal at other than the instability frequency. Stability can occur only if the controller can generate a signal at the frequency of the unstable oscillation. Thus, if pulsed control is stabilizing, the control is achieved through the harmonic of the control signal that corresponds in frequency with the unstable oscillation frequency. For the subharmonic case, the controller is effectively emulating a simple phase shift controller by first dividing the input frequency to generate a subharmonic pulse and then multiplying the frequency, via the harmonics of the pulsed waveform, to achieve an output at the input frequency.

This mechanism of controller operation means that pulsed control will be less efficient in terms of actuator power than linear control because energy in any harmonics other than that at the instability frequency will be wasted. Also, although the advantage of subharmonic forcing is reduced cycling of the actuator, the bandwidth of the actuator must be sufficient to generate significant energy at the oscillation frequency, and hence the actuator bandwidth must be on the order of the oscillation frequency. There is no real savings in bandwidth by using subharmonic control. The analysis that follows assumes the actuator is driven by a square pulse and the system responds to this pulse in a linear manner. For acoustic actuators, this is a reasonable assumption and thus we use acoustic actuators in our experimental verifications. For fuel modulation actuation, the input to the acoustic system is the heat release pulse that results from burning the fuel perturbation. The effects of diffusion, mixing, and flame shape will cause the heat release pulse to be a distorted version of the original fuel pulse. To the extent that these effects are linear, the following analysis will hold. Significant nonlinearities will introduce additional complexities that will need to be considered in future work.

5.2.2 Fixed Height Control Signals

The first case to be considered is one in which the output pulses of the controller have a fixed amplitude, X , and a fixed duration, w , as shown in Figure 5.3.

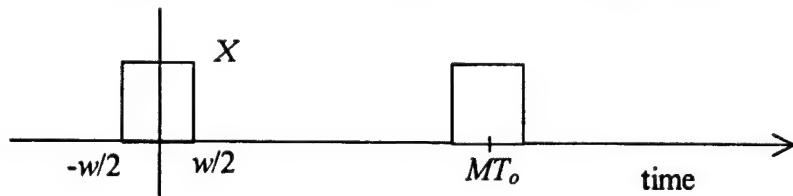


Figure 5.3 - Typical Pulse Train

For a sine wave input at frequency ω , amplitude A , and period T , the controller outputs a pulse train with on-time w , height X , and period MT . The Fourier expansion of this pulse train can be written as

$$\sum_{k=0}^{\infty} 2|C_k| \sin\left(\frac{k\omega}{M}(t - \tau - \frac{w}{2}) + \angle C_k + \frac{\pi}{2}\right),$$

where the coefficients C_k are given by

$$C_k = \frac{wX}{MT} \text{sinc}\left(\frac{wk\pi}{MT}\right).$$

The magnitude of the harmonic at the oscillation frequency ($k=M$) is

$$2C_M = \frac{2wX}{MT} \text{sinc}\left(\frac{wM\pi}{MT}\right) = \frac{2X}{M\pi} \sin\left(\frac{w}{T}\pi\right) \quad (5.1)$$

Two significant conclusions can be drawn from this formula. First, to maximize the energy at the oscillation frequency for a fixed amplitude pulse (essentially to maximize the efficiency of the controller), the on-time of the pulse should be $w=T/2$, or half the period of the frequency we are trying to control. Previous work with variable duty cycle signals [9] did not consider the relationship to the Fourier components of the signal. Second, as the order of the subharmonic, M , increases, the amplitude of the pulse must increase proportionally to produce the same effect at the oscillation frequency. Thus, as the cycle requirements of the actuator are reduced, the pulse amplitude produced by the actuator must increase and the on-time of the actuator should remain constant, if the effect of the actuation is to remain constant.

The above conclusions also make sense from an energy (Rayleigh criterion) viewpoint. From this view, the goal of control is to extract sufficient energy from the system so that oscillations cannot grow and are instead damped out. To accomplish this, the controller must provide heat release or positive acoustic forcing during the negative half cycle of the pressure perturbation. Thus, a pulse on-time of $T/2$ makes sense in this context. In addition, if you extract energy only every M^{th} cycle, then the amount of energy input by the actuator must increase by a factor of M for the net effect to remain the same.

The describing function of the controller is given by

$$DF = \frac{2X}{AM\pi} \left| \sin\left(\frac{w}{T}\pi\right) \right| \angle \left[\frac{\pi}{2} + \angle C_M - \omega\left(\tau + \frac{w}{2}\right) \right]. \quad (5.2)$$

Since the delay, τ , is a control variable, any desired phase can be realized with this controller. If the phase of the controller is fixed at a value that allows stabilization of the system, then questions about the required gain of the controller can be determined as in Figure 5.4.

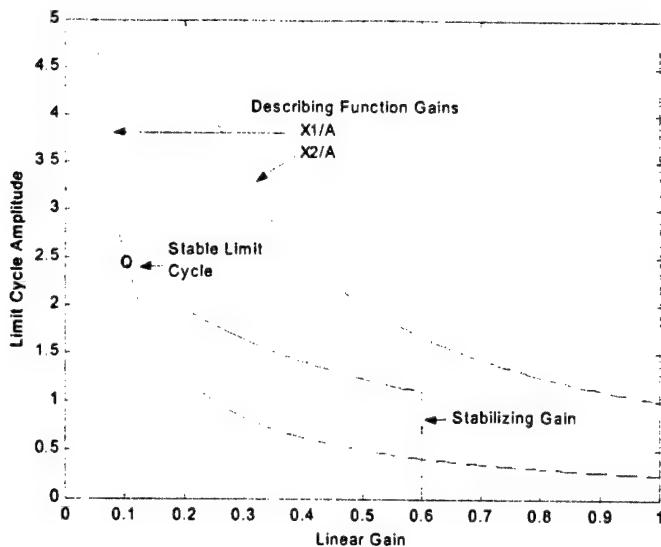


Figure 5.4 - Describing Function with Fixed Amplitude Pulse Heights

Assuming a linear (proportional actuation) phase-shift controller with the same phase as the pulsed controller, one can plot the limit cycle amplitude as a function of the linear controller gain from zero gain until the controller stabilizes the system. Then, plot the effective linear gain, which is the amplitude of the describing function, as a function of limit cycle amplitude. Since the describing function gain varies as a constant divided by the limit cycle amplitude, the gain will approach zero as the limit cycle amplitude approaches infinity, and the gain will approach infinity as the limit cycle amplitude approaches zero. Two cases for the describing function gain are shown in Figure 5.4, corresponding to two different values of the ratio X/A , where $X_2 > X_1$.

For the case of height X_1 , where the describing function gain intersects the limit cycle amplitude curve, there are two possible solutions. One is a limit cycle with amplitude corresponding to the intersection point and the other corresponds to stabilization of the system, which will be discussed in the next section. If the limit cycle is already established before the controller is turned on, it will have an amplitude corresponding to zero controller gain, which is the intersection of the limit cycle curve with the vertical axis. If the controller is switched on, the limit cycle will have its amplitude reduced to that corresponding to the intersection point on the graph. The system will not be stabilized. On the other hand, if the controller and the system are "turned on" simultaneously, the limit cycle will not develop because the gain of the controller exceeds the gain needed for stabilization. Note that this will be the case for any controller with fixed pulse amplitude, because the gain always tends to infinity for small input amplitudes. In practice, if the system is stable, a disturbance can occur that will increase the oscillation amplitude to a point where the effective control gain is no longer sufficient for stabilization. In this case, the system will jump to the stable limit cycle. Similarly, if the system is at the stable limit cycle point, a disturbance could reduce the amplitude sufficiently to cause a jump to the stabilized condition.

In the second case pictured in Figure 5.4, corresponding to height X_2 , there is no intersection between the effective gain and the limit cycle curve. The system will always be stabilized because for any limit cycle amplitude the control gain is greater than the gain needed to stay at that amplitude. Thus, the oscillation amplitude will be forced to decrease

with the eventual result that the control gain will become greater than the gain needed for stabilization. If the limit cycle amplitude curve were actually available for a system, it would be straightforward to compute values of w , X , and M to guarantee stabilization.

A third case, involving two intersections, is also possible and is pictured in Figure 5.5. One intersection corresponds to a stable limit cycle, meaning that nearby trajectories will be attracted to the limit cycle solution. The other intersection corresponds to an unstable limit cycle, meaning that nearby trajectories will be repelled from this limit cycle and stabilize elsewhere. In practice, an unstable limit cycle cannot be observed, since the smallest perturbation from this theoretical solution will cause the system to move away from this solution. If the oscillation amplitude is above the amplitude of the unstable limit cycle, the gain is such as to cause the system to approach the stable limit cycle. If the amplitude is below that of the unstable limit cycle, the system will be stabilized.

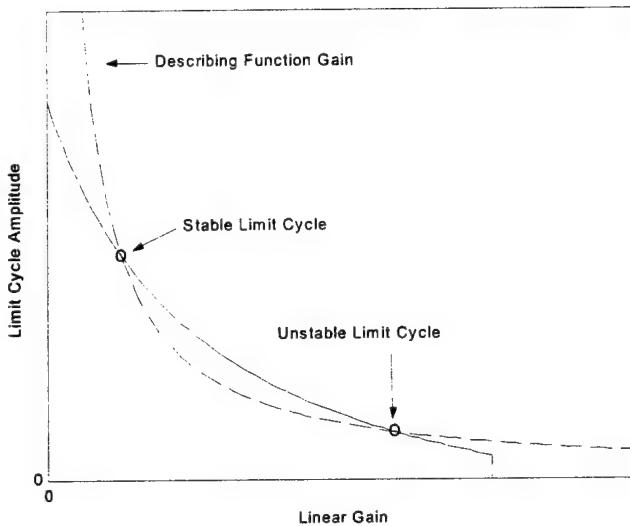


Figure 5.5 - Describing Function with Two Intersections

5.2.3 Ultimate Amplitude with Fixed Forcing

Although the discussion above assumed that if the original limit cycle was no longer viable the system would be stabilized, this is not strictly true. From a physical perspective, a fixed height and fixed duration pulse will impart a finite energy into the system. Since control action is needed to avoid unbounded growth and any control action imparts finite energy, a stabilized state having zero pressure oscillation is not possible. From a control point of view, once the effective gain is increased beyond the stabilizing gain, k_{stab} , the system will begin to move towards the zero oscillation state causing the effective gain of the controller to increase towards infinity. For any system with a pole-zero excess of three or more (which will include any real control system), increasing the control gain without bound will eventually cause one or more poles to go unstable at an ultimate gain denoted by k_{ult} . This situation is pictured in the complex plane in Figure 5.6.

When the effective gain increases beyond k_{ult} and a pole becomes unstable, the pressure oscillation will begin to increase again, causing the effective gain to decrease. This

process will result in a stable limit cycle such that the effective gain is equal to k_{ulr} , leading to the equation

$$\frac{2X}{AM\pi} = k_{ulr} \text{ or } A = \frac{2X}{k_{ulr} M\pi} \quad (5.3)$$

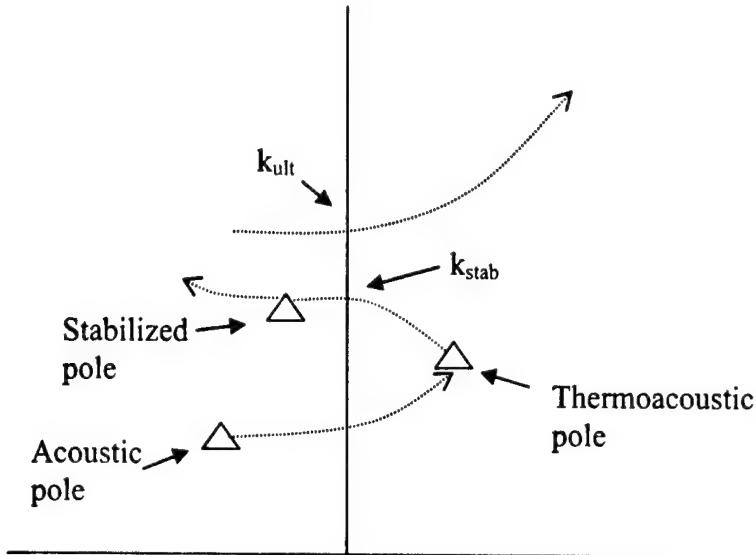


Figure 5.6 - Depiction of Pole Movement in the Complex Plane as the Control Gain is Increased

The latter equation shows that the amplitude of the ultimate limit cycle will be proportional to the pulse height and inversely proportional to the harmonic number M and the ultimate gain k_{ulr} . Thus, although the system is never stabilized to zero, it achieves a stable limit cycle caused by the pulsed nature of the controller. The amplitude of this limit cycle can be controlled by the choice of M and perhaps by using a more sophisticated control design in place of the phase-shifting delay so as to increase the value of k_{ulr} . In Figures 5.4 and 5.5, this limit cycle can be pictured as occurring at the intersection of the describing function curves with a vertical line drawn at k_{ulr} .

Thus, we see that there are two limit cycles to be considered. The first is due to the thermoacoustic interaction, with its amplitude limited by the nonlinear characteristics of the heat-release dynamics. The second is due to a high-gain controller, with its amplitude limited by the nonlinear gain characteristic of the fixed-amplitude control signal. For maximum effectiveness in eliminating the thermoacoustic limit cycle, it is desirable to maximize control gain by using $M=1$. After the system has transitioned to the control-induced limit cycle, it is desirable to minimize the resulting amplitude by operating at the highest practical value of M . This leads us to propose the idea of a variable subharmonic controller, where $M=1$ is used for initial stabilization and a higher value is used to minimize the amplitude of the control-induced limit cycle and reduce cycle fatigue of the actuator.

5.2.4 Proportional Control Signals

If the intention is to use pulsed control to drive the system oscillations to zero, then the amplitude of the output pulses of the controller must be proportional to the controller

input. This proportionality will enable the effective gain of the controller to be set between k_{stab} and k_{unst} ensuring stability of the system. In the ideal case, a proportional control would be represented by a vertical line on the gain-amplitude plots we have been considering. In reality, there will be some maximum pulse amplitude for the actuator, after which it will saturate, essentially transitioning to a fixed-height signal for suitably large inputs. This situation is shown in Figure 5.7. Clearly, the three previous cases we have discussed – zero, one and two intersections with the limit cycle amplitude curve are possible in the case of a saturating, proportional actuator. In addition, a case with three intersections is also possible and is pictured in Figure 5.7. This case exhibits a proportional gain that is less than the required stabilizing gain. The result is two possible stable limit cycles separated by an unstable limit cycle.

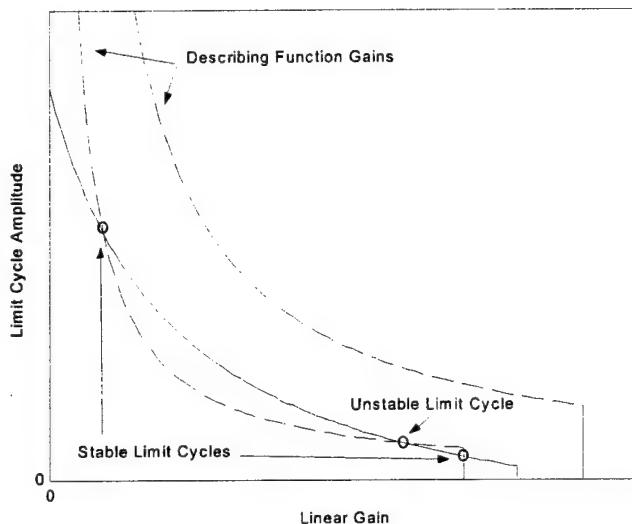


Figure 5.7 – Describing Function with Proportional Height Pulses

Depending on the nonlinearity of the heat release dynamics, the limit cycle amplitude curves could exhibit much more complicated behavior than those illustrated in the curves of the above figures. Thus, cases other than those outlined above are possible, but the previous analysis should serve to understand these new cases as well.

5.3 Controller Implementation

5.3.1 Terminology

In this chapter a subharmonic signal will be referred to as a $1/M x\%$ signal, where M is the order of the subharmonic and x is the duty cycle of the control signal. Therefore a $1/3$ signal will be a subharmonic signal with a pulse every third positive-going zero crossing of the pressure signal. Note that the duty cycle is the percentage of time that the *control* signal is high. Therefore a $1/2$ 25% signal will have the same pulse width as a $1/1$ 50% signal, but will occur only every other period of the limit cycling pressure signal. Examples of $1/1$ and $1/2$ signals with equal pulse widths are shown in Figure 5.8.

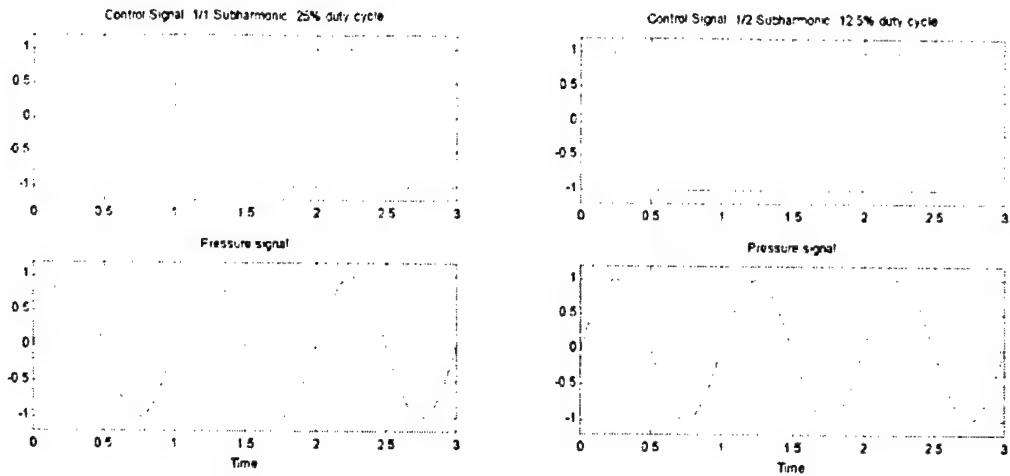


Figure 5.8 - Typical Pulsed Control Signals

5.3.2 Explanation of Control Algorithm

The algorithm was written in C and implemented on a DSpace DS1103 controller. The algorithm generates a pulse train in the time domain, which is based on detecting zero crossings of the input signal, in this case the pressure fluctuation.

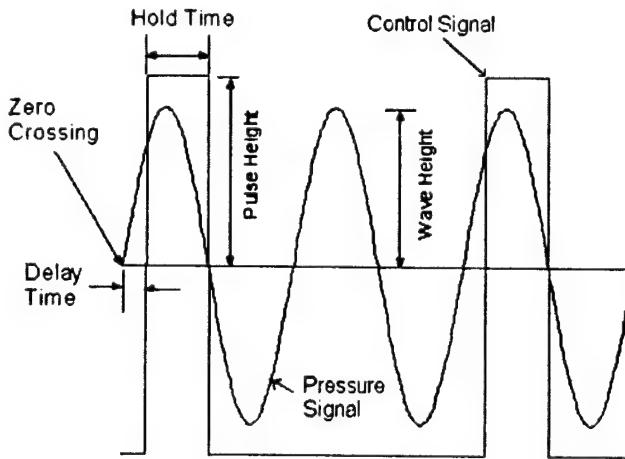


Figure 5.9 - Control Algorithm Illustration

As shown in Figure 5.9, the algorithm starts a timer when it sees the previous sample of the input less than zero and the current sample greater than zero. When this timer reaches the value "Delay Time," a user selectable number of samples, it will set the control signal high. The height of the pulse can be fixed or made proportional to the wave height with a user-selectable gain. When the pulse is initiated, the hold timer is started. When this timer gets to the "Hold Time" number of samples, the pulse is ended and the signal goes to the negative value of the height calculated previously. The algorithm then returns to detecting zero crossings. When the set number of positive-going zero crossings M is reached ($M=2$ in Figure 5.9) the delay timer is started again and the next pulse is output.

5.4 Experimental Results

5.4.1 Experimental Setup

A block diagram of the experimental system is shown in Figure 5.10. The combustor is a tube that is acoustically closed at the bottom and open at the top. Premixed methane and air are injected at the bottom and a ceramic honeycomb flame holder is located at the tube midpoint. The second acoustic mode of the tube, at a frequency near 180Hz, goes unstable for a wide range of equivalence ratios and flow rates [9]. The pressure transducer signal in the combustor is filtered before the A/D conversion to eliminate spurious zero crossings. The D/A output is low-pass filtered to protect the speaker, as a speaker is not mechanically designed for pulse inputs. However, all harmonics up to and including the instability frequency are allowed to pass into the combustion system. The signal is sampled at 10 kHz to allow for as great a resolution as possible in setting the phase and duty cycle.

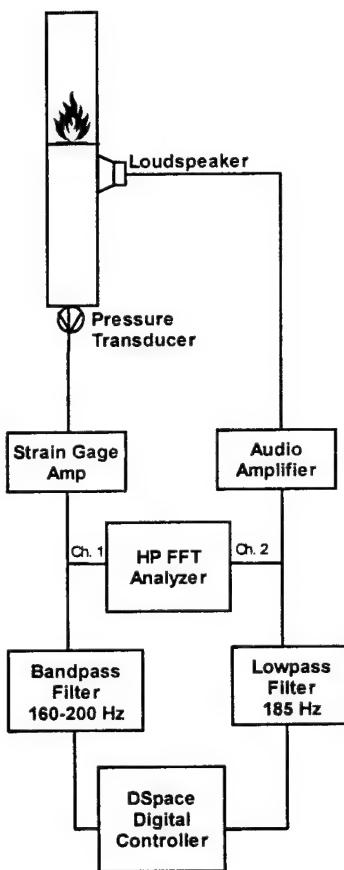


Figure 5.10 - Test System Layout

5.4.2 Linear Phase Shifter Results

5.4.2.1 Hysteresis Curve

The linear phase shifter simply outputs the input signal after a delay, thus passing all frequency components of the input (within the sampling rate and filtering constraints). Experiments were first performed using a fixed gain linear phase shifter to determine the optimal phase shift to be used. This phase shift was then used for all subsequent phase shift

experiments. It was found that a delay of $\sim 288^\circ$ (45 samples at 10 kHz sampling rate) produced the best results and the lowest gain for stabilization.

The amplitude of the limit cycle is shown as a function of the linear control gain in Figure 5.11. It was found that it took less gain to maintain stability than to achieve it. Control is achieved at a gain of 1.34 and lost when the gain drops below 1.16.

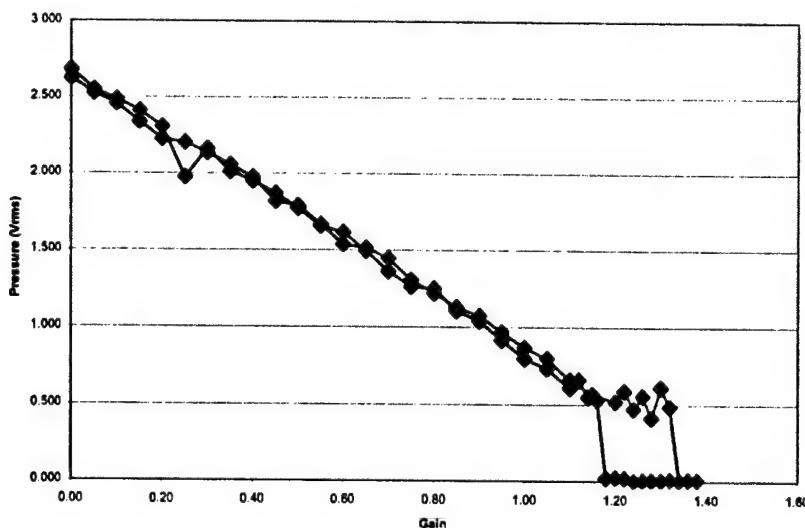


Figure 5.11 - Linear Phase Shifter Hysteresis Curve

5.4.2.2 Loss of Control at High Gains – Peak Splitting

If the gain of the linear phase shifter is increased past the gain needed for stabilization, a secondary instability results. Gains above 2.8 resulted in a power spectrum with a split peak, indicating two potential instabilities evenly spaced around the original instability frequency. Experimentally it was determined that a gain of 4.2 would cause a secondary instability. This instability results in a new limit cycle that is at a different frequency than the original limit cycle. Without an accurate mathematical model of the system it is not possible to know whether this new limit cycle is due to the acoustic pole, which was stabilized, moving back to instability, or whether one of the bandpass filter poles in this same frequency band is becoming unstable. Practically, it doesn't matter.

The peak splitting is because two different poles are going unstable at nearly the same time. The phenomenon can be seen in Figure 5.12Figure as the gain is increased in increments of 1.5 from 0 to 4.5. Initially the thermoacoustic limit cycle exists. As the gain is increased to 1.5, control is achieved, leaving a lightly-damped, noise-excited acoustic peak. As the gain is further increased to 3.0, peak splitting occurs. The acoustic peak at the instability frequency is visible, but has been driven further down and two additional peaks have arisen to either side of it. Finally, at the gain of 4.5, stability is again lost. This characteristic has been well documented for phase shift controllers [9]. The characteristic is presented here to fully characterize the linear phase shifter's effects on the plant.

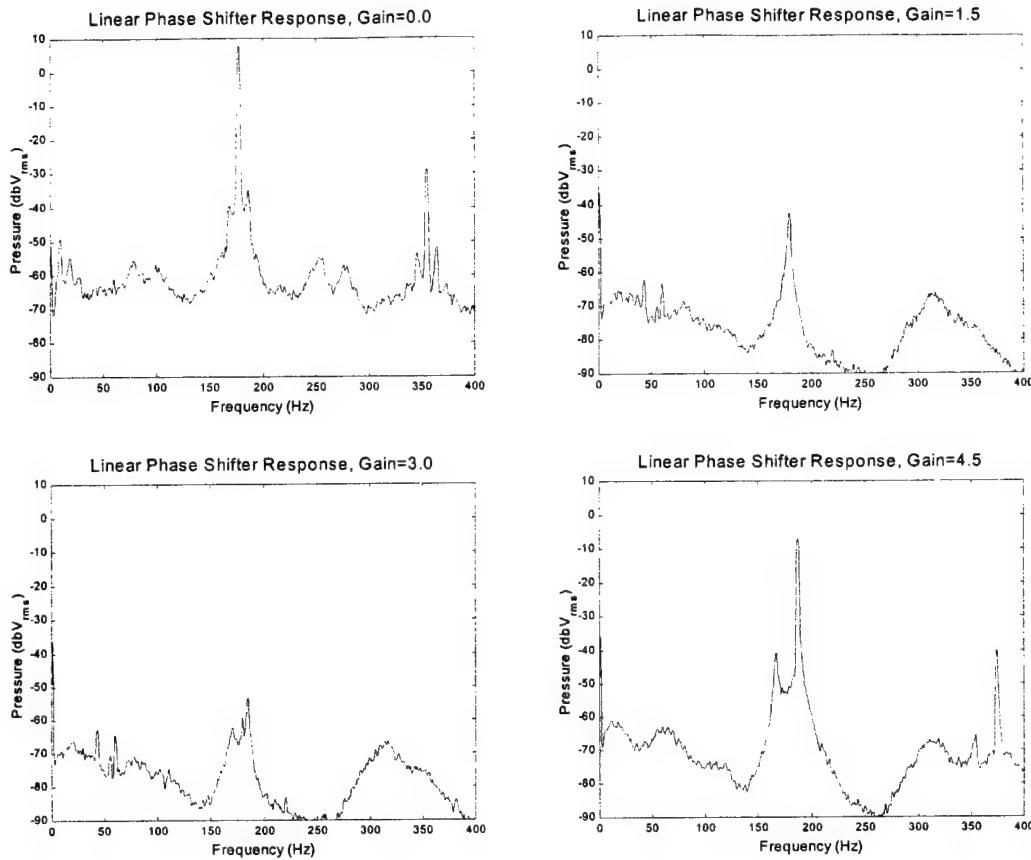


Figure 5.12 - Control Level with Linear Phase Shifter, Various Gains

5.4.3 Validation of the Instability Frequency Hypothesis for Pulsed Control

If the mechanism of control in the tube combustor is indeed the component of the control signal at the instability frequency, we can support this by examining the ratios of the gains needed to achieve stability for pulse train control signals of various duty cycles and subharmonic ratios. In the experiments, a pulsed-control signal is generated with a specified duty cycle and subharmonic ratio that has an amplitude related to the input signal amplitude by a gain g . The gain is increased until the limit cycle is stabilized. The value of the gain at stabilization is recorded. The effective control signal gain at the instability frequency is gF , where F is the Fourier coefficient associated with the frequency component at the instability frequency (i.e. the M th coefficient) for a unit-amplitude pulse train having the specified duty cycle. If control is obtained due to the component of the control signal at the instability frequency, then the observed gains from two experiments using different duty cycles or subharmonic ratios should be related by $g_1F_1 = g_2F_2$. This implies that the ratio of the gains should equal the inverse of the ratio of the Fourier coefficients.

5.4.3.1 One-Half Subharmonic Control

For the 1/2 signal case, experiments were done at a 25% duty cycle, 17% duty cycle, and 10% duty cycle. Experiments also were done with a linear phase shifter, and a 1/1 signal with duty cycles of 20%, 34%, and 50%. These duty cycles give pulse widths equal to the 1/2 signal cases. All tests were done with the same phase shift of the control signal as determined by the HP analyzer, which monitors the total phase of the control path including filters, as seen in Figure 5.10.

Experimentally it was determined that the gain required to achieve control was 1.90 for the 1/2 25% case and 1.76 with a 1/1 20% case, giving a ratio of 0.926. The second Fourier coefficient for a 25% signal is 0.6366, and the first Fourier coefficient of a 20% signal is 0.7354, producing an inverse ratio of 0.866. Therefore, there is a 7.0% error between the observed gain and the expected gain based on the hypothesis. The actual ratios of the gain of every signal type tested to the gain of every other signal were computed and are shown in Figure 5.13 as a function of the expected ratios of the gains. The gains for achieving and losing control were both measured and both corresponding ratios are shown. These will be different due to the hysteresis of the system. Ideally the points should all lie exactly on the 45° line shown, with the expected ratio equal to the actual ratio. The subharmonic signals agree very closely when compared to other subharmonic signals, and the fundamental signals agree very closely when compared to other fundamental signals and the linear phase shifter. However, there is a slight amount of error between the ratios of subharmonic signals and the fundamental, as well as the subharmonic signals and the phase shifter. This is most likely due to non-linear effects in the combustor where it has been observed that a low-frequency signal, not harmonically related to the limit cycle, injected with a control signal can act to reduce the gain necessary to control the system. In general, however, the test results agree closely with the expected results of the linear analysis and appear to verify the hypothesis that the limit cycle frequency component of the control signal is responsible for control of the system and not some nonlinear interaction of the subharmonic frequency component with the system.

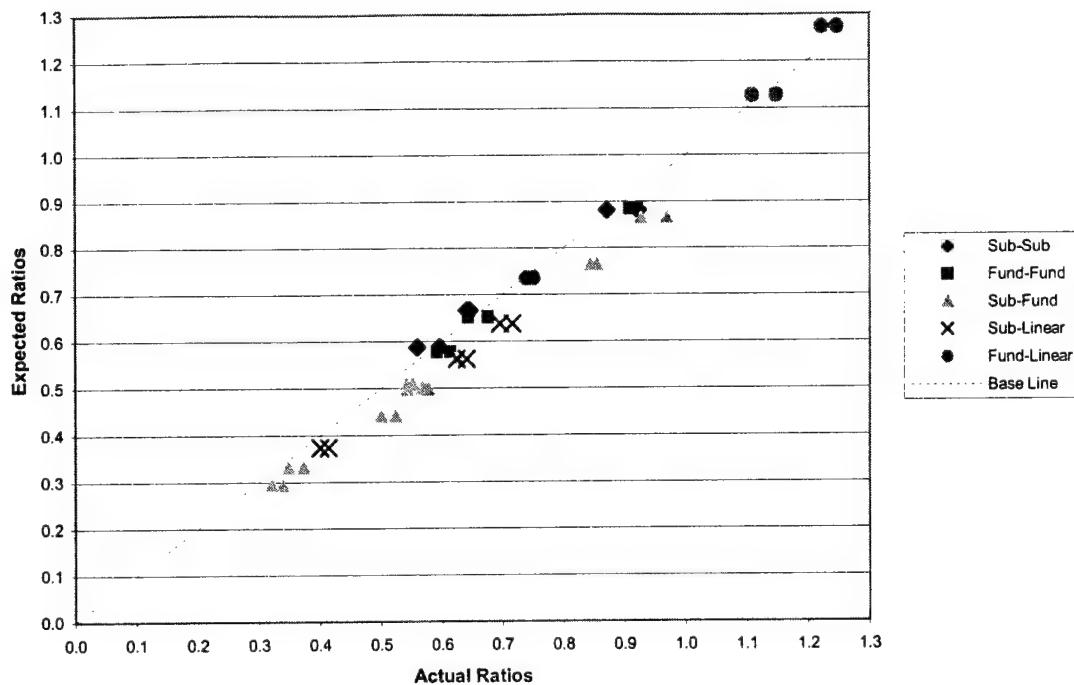


Figure 5.13 - Actual vs. Expected Ratios for Half-Harmonic Signals

5.4.3.2 One-Third Subharmonic Control

Similar tests were done using a one-third subharmonic control signal. Duty cycles of 8%, 11%, and 16% were tested. These duty cycles have pulse widths equal to those considered in the previous section. As above, ratios between every case were calculated and compared to theoretical expectations. The actual and expected ratios are plotted in Figure 5.14. The ratios between gains at different duty cycles for control using both fundamental frequency control and subharmonic control match very nearly with what is expected. Ratios between subharmonic and fundamental pulse trains and ratios compared to a linear phase shifter are somewhat farther away from the expected than in the half-harmonic case. Having two subharmonic frequency components present seems to enhance the nonlinear effect noted in 5.4.3.1, requiring even less gain to control the system. Overall, however, the results are close to expected and appear to verify that the Fourier component at the limit cycling frequency is the dominant means of control.

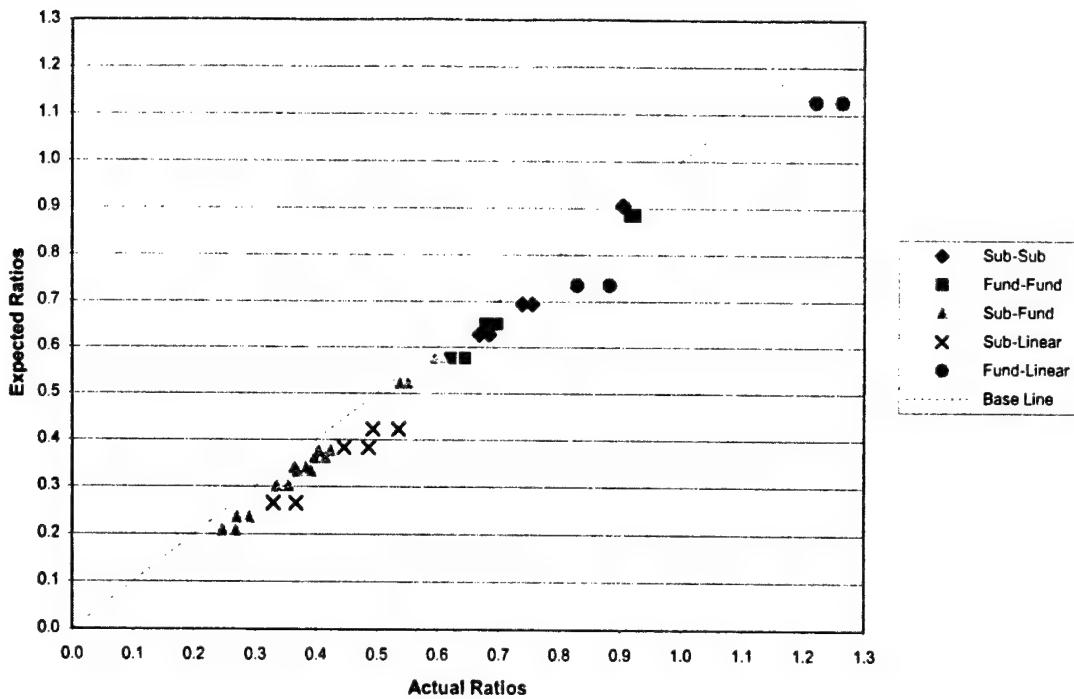


Figure 5.14 - Actual vs. Expected Ratios for Third-Harmonic Signals

5.4.4 Validation of Fixed Pulse Height Analysis

According to the theory in Section 5.2.2, the pulse height required to stabilize the system can be calculated from the describing function amplitude given in (2) and the system response with a linear phase shifter shown in Figure 5.11. For a 1/1 50% signal ($M=1$, $w/T=0.5$) the describing functions for various pulse heights are plotted and superimposed on the hysteresis curve in Figure 5.15. As can be seen in the plot, pulse heights above 0.8 V should result in stabilization of the thermoacoustic limit cycle, as there are no intersections for pulse heights above 0.8 V. The actual value of the limit cycle with each different control signal is shown as a line in Figure 5.15, and it can be clearly seen that the actual level is at the intersection of the describing function and the linear phase shifter results.

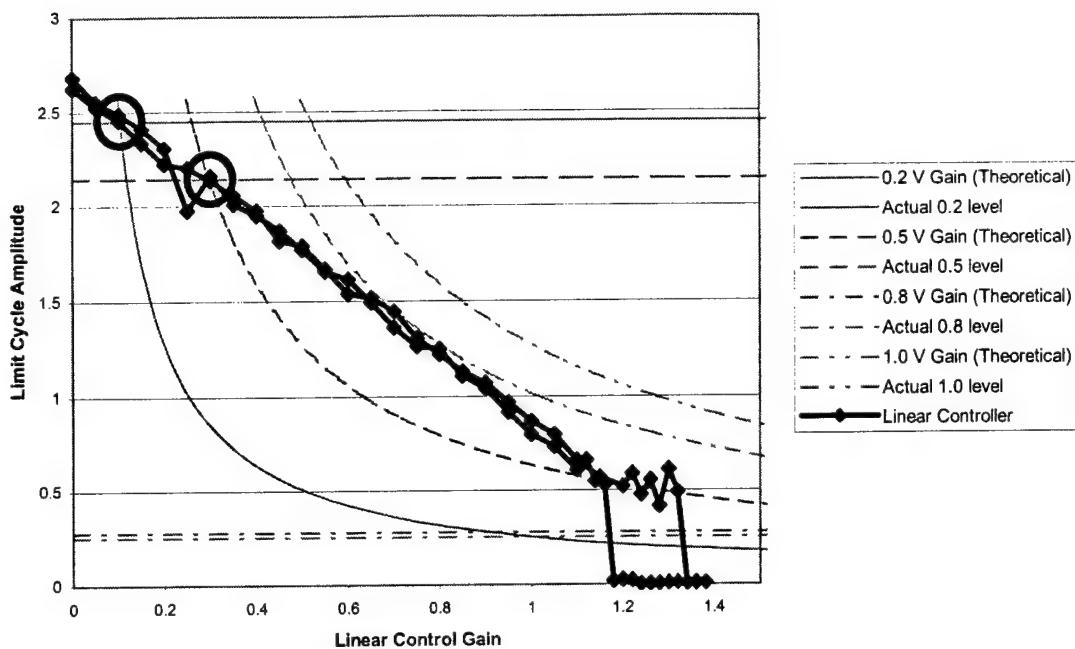


Figure 5.15 - Describing Function Plots Superimposed on Hysteresis Curve

Not pictured in Figure 5.15 are the control-induced limit cycles for each condition, which exist further to the right where the linear gain is 4.2. Using any control gain where both intersections exist (thermoacoustic and control-induced), it was fairly easy to move between the thermoacoustic and control-induced limit cycles. A disturbance created by interrupting the airflow at the top of the combustor tube was enough to cause the combustor to jump from one instability to the other, with no change to the controller. The farther apart the two limit cycles, the greater the disturbance needed to cause this jump. As the pulse height increased and the two limit cycles became closer in amplitude, the jump would occur without providing an external disturbance. In these cases the control-induced limit cycle was the "dominant" limit cycle.

A plot of the limit cycle amplitude as a function of pulse height is given in Figure 5.16. This plot shows that the transition from the thermoacoustic limit cycle to the control-induced limit cycle takes place above an amplitude of 0.5 V. Once the system has jumped to the control-induced limit cycle, it is possible to reduce the pulse height and still remain on the control-induced limit cycle, taking the lower branch of the figure. For pulse heights between 0.6 V and 0.8 V an external disturbance would cause the system to temporarily jump to a thermoacoustic limit cycle and then immediately return to the control-induced limit cycle. This explains the discrepancy between the break-off point at 0.6 V as shown in Figure 5.16 and the expected break-off point of 0.8 V shown in Figure 5.15.

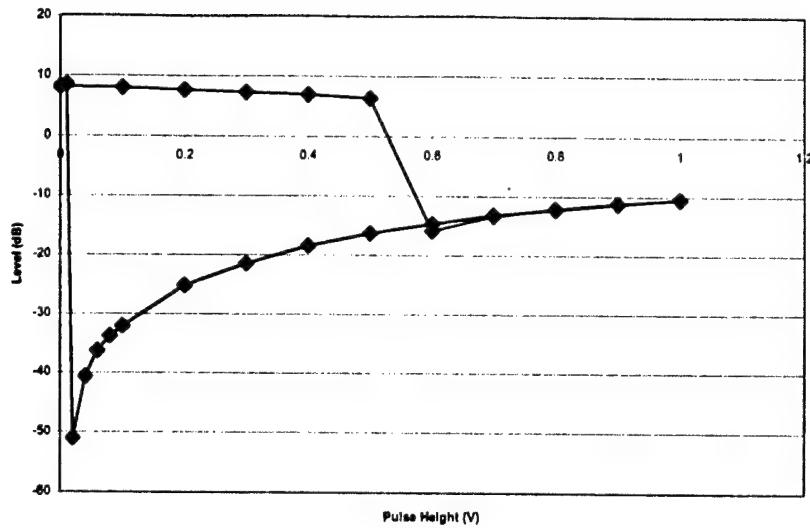


Figure 5.16 - Instability Level for Fixed Pulse Height Controller

The control-induced limit cycles should correspond to a gain of $k_{ul}=4.2$. Examining the control-induced limit cycle with a pulse height of 0.9 V, for instance, it is seen that the maximum level of the limit cycle is 0.276 V. The gain (multiplied by 2 because a ± 0.9 V signal was actually used) can then be found to be

$$2 \frac{2(0.9)}{0.276\pi} = 4.15$$

The predicted values of k_{ul} for the control-induced limit cycles at all pulse heights are shown in Figure 5.17. It can be seen that the gain tracks very closely with the expected value of 4.2. The inaccuracy at the 0.08 V point is likely due to the extremely small amplitude (9 mV) of the instability and the possibility of noise at this low level.

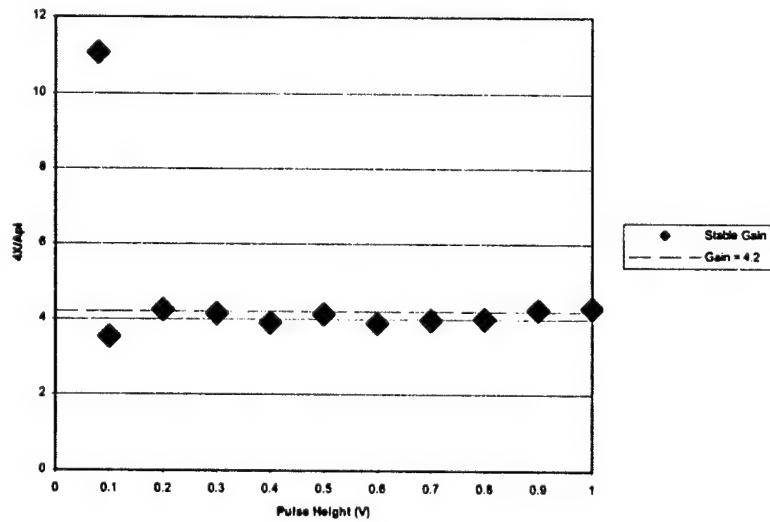


Figure 5.17 - Gain of Fixed Pulse Height Controller

Experiments were also done with variable duty cycle, fixed height control signals, using a subharmonic ratio of one. Note from equation (1) that the Fourier component of the control signal is periodic, with the maximum value occurring at a duty cycle of 50%. This means the level of control with a fixed height controller should be maximized with a duty cycle of 50%, and fall off symmetrically with changes in duty cycle away from 50%. Thus a 40% duty cycle signal and a 60% duty cycle signal should have exactly the same results, since the Fourier component at the instability frequency is exactly the same in both cases.

Experiments were done to verify this. A fixed height signal was input to the system at 40%, 50%, and 60% duty cycles. The phase shift of the controller was varied until the frequency of the instability was 186 Hz in each case. As can be seen in Table , the assumption that the control amplitude is maximized at 50% and that the level is identical at 40% and 60% is verified. The level of suppression goes down at 50%. The increased amplitude of the control signal at 50% duty cycle causes the amplitude of the ultimate limit cycle to be larger, in accord with the argument given in Section 5.2.3.

Duty Cycle	Pressure (dbVrms)	Phase Shift (°)	Frequency (Hz)
40%	-6.98	65.7	186
50%	-5.83	65.7	186
60%	-6.96	65.5	186

Table 5.1 - Level of Control with Variable Duty Cycles

5.5 Conclusion

This chapter has analyzed the effects of pulsed control, both at the fundamental (limit cycling) frequency and subharmonic frequencies, on the stabilization of thermo-acoustic instabilities. The hypothesis that the control effectiveness depends on the amplitude of the control signal at the instability (input) frequency was argued theoretically and verified experimentally using acoustic actuators. Methods for determining the pulse height necessary to eliminate the thermo-acoustic limit cycle and for calculating the amplitude of the resulting control-induced limit cycle were derived for fixed-pulse-height systems and verified experimentally. Finally, based on this theory, a variable-subharmonic control scheme was proposed that uses fundamental forcing to transition the system from the thermoacoustic limit cycle to the control-induced limit cycle, and then uses subharmonic forcing to reduce the level of the control-induced limit cycle, while also reducing the cycles required of the actuator.

Although we expect the basic ideas presented in this chapter to be useful for understanding the behavior in more complex, practical combustors, it is understood that additional complexities arise, such as flame repositioning, that create added challenges to effective control.

5.6 References

- [1] Annaswamy, A.M., Fleifil, M., Rumsey, J. W., Prasanth. R., Hathout, J.P., and Ghoniem, A.F., 2000, "Thermoacoustic Instability: Model-based Optimal Control Designs and Experimental Validation," *IEEE Transactions on Control Systems Technology*, **8**, Issue 6, pg. 905-918.
- [2] Kim, K., Lee, J., Stenzler, J., and Santavicca, D., 2000, "Optimization of Active Control Systems for Suppressing Combustion Dynamics," RTO Symposium on Active Control Technology for Enhanced Performance Operational Capabilities of Military Aircraft, Land Vehicles and Sea Vehicles, Braunschweig, Germany, May 8-11.
- [3] Jones, C M., Lee, J.G., and Santavicca, D.A., 1999, "Closed-loop active control of combustion instabilities using subharmonic secondary fuel injection," *Journal of Propulsion & Power*, **15** n 4, p 584-590.
- [4] Heising, R., Lubarsky, E., Neumaier, M., Neumeier, Y., and Zinn, B.T., 2000, "Periodic Liquid Fuel Sprays Combustion Processes and Their Damping of Combustion Instabilities," AIAA 00-1024, 38th AIAA Aerospace Sciences Meeting, Reno, NV, Jan. 10-13.
- [5] McManus, Keith R., Magill, John C., and Miller, Michael F., 1998, "Combustion Instability Suppression in Liquid-Fueled Combustors," AIAA 98-0642, 36th AIAA Aerospace Sciences Meeting, Reno, NV, Jan. 12-15.
- [6] Fleifil, M., Hathout, J. P., Annaswamy, A. M., and Ghoneim, A. F., 2000, "Reduced Order Modeling of Heat Release Dynamics and Active Control of Time-Delay Instability," AIAA 00-0708, 38th AIAA Aerospace Sciences Meeting, Reno, NV, Jan. 10-13.
- [7] Hantschk, C., Hermann, J., and Vortmeyer, D., 1996, "Active instability control with direct-drive servo valves in liquid-fueled combustion systems," International Symposium on Combustion, **2**, Combustion Inst, Pittsburg, PA, USA. p 2835-2841.
- [8] Ogata, K, *Modern Control Engineering*, Prentice Hall, 1990.
- [9] Saunders, W.R., Vaudrey, M.A, and Eisenhower, B., 1999, "Perspectives on Linear Compensator Designs for Active Control," AIAA 99-0717, 37th AIAA Aerospace Sciences Meeting, Reno, NV, Jan. 11-14.

6 Conclusions

This work has involved the investigation, implementation and testing of three different types of adaptive algorithms for the suppression of thermoacoustic instabilities and optimization of combustor performance. The algorithms have different strengths and weaknesses, but all have proven effective in suppressing thermoacoustic oscillations, and two of the algorithms will clearly be useful for the slower task of optimizing combustor performance.

The pattern search algorithms are the slowest, yet most robust of the algorithms tested. Since they make no assumptions about the performance surface, they are less likely to become trapped in local minima and can traverse rough performance surfaces. Thus, pattern search algorithms are ideal for optimization of combustor performance. What was perhaps surprising was how well they performed at the more time critical task of stabilizing thermoacoustic instabilities. As the order of the adapted filters becomes higher, or the combustor becomes noisier (more turbulent), it is expected that this class of algorithm might become too slow for the stabilization task.

The explicit gradient algorithms are faster than the pattern search algorithms because they assume more smoothness in the performance surface and attempt to move in the best downhill direction. For rough performance surfaces, however, this could cause problems. In the case of thermoacoustic oscillations, these algorithms proved very effective in our experiments, both with linear and on-off actuation. For noisier combustors, this class of algorithm will also slow down due to the increased integration time needed. On the positive side, both explicit gradient and pattern search algorithms require very little *a priori* information about the process being controlled, which makes them quite robust. In addition, their digital implementation makes it easy to incorporate constraints so that these algorithms will never remain in a state that makes the performance worse than the best previous performance achieved.

The filtered-E algorithm was developed to avoid an explicit computation of the gradient. The gradient is computed using information from a linearized model of the control to error path. This makes the algorithm much faster than the other algorithms considered, as its speed does not depend on an integration interval or scale linearly with the number of parameters adapted. The algorithm requires more *a priori* information, which must be acquired in a real time system identification step prior to starting the algorithm. This can limit the robustness of the algorithm and probably makes it unsuitable for working with the rough performance surfaces that could be encountered with optimizing combustor performance metrics, such as efficiency. On the other hand, this algorithm should work well in the noisy environments characteristic of highly turbulent combustors. It was very effective in suppressing thermoacoustic instabilities in our experiments.

Our work on pulsed control of combustors has shown that there appears to be little benefit to using subharmonic actuation in terms of reducing the bandwidth required of the actuator. On the other hand, the reduced cycling of the actuator may be a major benefit. The variable subharmonic control that we have proposed should be useful not only in reducing actuator cycles, but also in increasing the achievable suppression when using on-off actuators.

Although we have achieved success on the simple combustors we have used for experiments, future work must concentrate on validating our techniques on higher power

liquid-fuel combustors. Some results for the pulsed control of a multi-injector liquid-fuel combustor have been obtained using the time-averaged-gradient algorithm, but much more experimental testing needs to be done using both proportional and on-off actuators and all of our algorithms. In addition, our pulsed-control analysis for subharmonic actuation must be validated using liquid fuel actuation.